



Ms Jo Briskey MP
Chair
Joint Select Committee on Artificial Intelligence
PO Box 6100
Parliament House
Canberra ACT 2600

14 September 2026

Dear Chair,

Submission to Joint Select Committee on Artificial Intelligence

Thank you for the opportunity to provide a submission to the Joint Select Committee on Artificial Intelligence.

Artificial intelligence (**AI**) presents a significant opportunity to lift Australia's productivity, strengthen economic resilience and create new sources of growth. The emerging challenge is increasingly one of value conversion. As Treasury has recently highlighted,¹ AI is the first substantive, credible, global productivity accelerant in almost two decades. Yet the gap between experimentation and transformation remains significant. Around two-thirds of medium to large Australian businesses report some AI use, fewer than one in ten describe that adoption as significant or at scale.²

Our experience with Australian organisations points to the same conclusion: value comes not from merely adopting tools, but from redesigning work. The focus must be on the transition from pilots to scaled

¹ Kehoe, J., 2026. AI the 'biggest economic transformation in our lifetime': Chalmers. Australian Financial Review, 31 August. Available at: <https://www.afr.com/policy/economy/ai-the-biggest-economic-transformation-in-our-lifetime-chalmers-20260831-p60svg> (accessed 11 September 2026).

² Reserve Bank of Australia., 2025. Technology investment and AI: What are firms telling us? RBA Bulletin, November. Available at: <https://www.rba.gov.au/publications/bulletin/2025/nov/technology-investment-and-ai-what-are-firms-telling-us.html> (accessed 11 September 2026).

PricewaterhouseCoopers, ABN 52 780 433 757
One International Towers Sydney, Watermans Quay, BARANGAROO
NSW 2000, GPO BOX 2650 SYDNEY NSW 2001
T: +61 2 8266 0000, F: +61 2 8266 9999, www.pwc.com.au
Level 11, 1PSQ, 169 Macquarie Street, PARRAMATTA NSW 2150, PO Box
1155 PARRAMATTA NSW 2124
T: +61 2 9659 2476, F: +61 2 8266 9999, www.pwc.com.au

deployment, from individual use cases to transformed workflows, and from technical capability to measurable outcomes.

AI's potential extends beyond efficiency. It can enable new products, services, business models and industries, particularly in sectors where Australia already has strengths, including resources, agriculture, healthcare, financial services, education and advanced manufacturing. The challenge is ensuring Australia captures more of that value domestically. This opportunity extends to government as well as business. AI has significant potential to improve public services, but Australia risks a two-speed transition if public sector adoption falls behind the broader economy.

The opportunity is productivity and growth. Not just from doing existing work smarter and faster, but doing different work altogether. In our view, four conditions will decide whether Australia seizes this opportunity:

- **Adoption has to reach scale:** helping organisations move beyond pilots to enterprise-wide implementation
- **Scale depends on AI infrastructure and sovereign capability:** securing investment in physical infrastructure, while strengthening Australia's domestic standards, testing and assurance capability
- **Investing in people provides the return:** managing the transition as AI reshapes jobs and entry-level pathways, while building the technical and human capabilities needed for higher-quality work and ensuring the benefits of AI are broadly shared across the workforce
- **Trust is the foundation on which everything else rests:** providing the governance, assurance and regulatory certainty needed to scale AI responsibly and safely.

While there exists optimism about AI's potential to improve productivity and deliver economic benefits, public sentiment is increasingly cautious. A 2026 Lowy Institute Poll found a large majority of Australians (64%) say the risks of AI outweigh the benefits, a 12-point increase since 2024.³ Many Australians have legitimate concerns around privacy, misinformation, workforce disruption, unequal access and more. Maintaining community confidence will require a balanced approach that supports adoption at scale while ensuring workers have access to the skills and opportunities needed to participate in an AI-enabled economy.


Our recommendations focus as much as possible on practical measures to accelerate adoption and value creation across business, government and the broader community in a responsible way: measuring

³ Lowy Institute., 2026. Artificial intelligence: Majority say risks outweigh benefits. 2026 Lowy Institute Poll. Available at: <https://www.lowyinstitute.org/publications/2026-lowy-institute-poll/artificial-intelligence-majority-say-risks-outweigh-benefits> (accessed 11 September 2026).

outcomes rather than usage, reducing barriers to adoption, strengthening assurance capability, and addressing targeted gaps in existing laws.

We look forward to further engagement with the Committee and would be pleased to discuss our submission with you. If you have any further questions, please do not hesitate to contact me.

Yours faithfully,



Kevin Burrowes
Chief Executive Officer

Realising Australia's AI opportunity

1. Productivity and growth

Term of Reference (a): the opportunities presented by AI to lift Australia's resilience, productivity growth, economic competitiveness, and living standards

Public discussion of AI's economic potential has focused largely on efficiency. AI can already help organisations perform existing work faster, at lower cost, and with greater quality by reducing administrative effort, accelerating analysis, improving service productivity and reducing repetitive work. These efficiency gains matter; productivity growth depends on creating more value from the resources available to the economy.

However, AI's most significant economic opportunity extends beyond improving existing operations. PwC's AI Performance Study⁴ found that AI delivers the most economic value when organisations use it to develop new products, services and ways of working – not simply improve what they already do. This includes making previously uneconomic services viable or combining data and expertise in new ways.

Australian organisations still remain focused primarily on efficiency and risk management. Fewer are focused on growing revenue, rethinking how they do business or pursuing new opportunities with organisations in other sectors. According to PwC's AI Performance Study, only 3% of Australian businesses reported using AI to improve business model transformation, compared with 59% of the study's top – performing companies globally, known as “AI leaders”. Similarly, only 10% reported using AI to support collaboration with organisations in other sectors, compared with 50% of AI leaders. These findings suggest Australia is yet to capture AI's broader potential for reinvention and growth.

Australia can build globally competitive AI-enabled products, services and intellectual property by combining its distinctive data, domain expertise and research capability in sectors where it has established strengths and significant unmet needs – including resources and energy, agriculture, health and care, financial services, education and advanced manufacturing.

Recommendation 1: The economic contribution of AI should be assessed by reference to the value it creates across the economy (including higher revenue, greater productivity, improved margins, increased capacity, better quality or a stronger competitive position), not simply by how widely it is used or how many use cases are deployed.

⁴ PwC Australia., 2026. PwC's AI performance study: Australian companies lead on AI security, lag on ROI. Available at: <https://www.pwc.com.au/media/2026/pwc-ai-performance-study-australian-companies-lead-on-ai-security.html> (accessed 10 September 2026).

2. Adoption at Scale

Term of Reference (d) the rate and extent of AI adoption across the Australian economy, and the barriers to adoption faced by small and medium businesses and family businesses

There is an important difference between giving people access to AI as individual users and using it at scale. Individual tools can improve productivity and build familiarity, but organisation-wide value requires reliable data and technology, redesigned processes, skilled people, clear accountability and practical ways to measure results.

PwC's AI Performance Study suggests that Australian organisations are willing to experiment with AI but are less successful at translating experimentation into performance. Just 13–14% report AI use that is scaled or embedded, compared with 41–44% of the study's AI leaders. Only 27% of Australian businesses reported that employees could quickly find and use high-quality data for current and future AI work, compared with 57% of AI leaders globally.

Scaling AI is a systems challenge: technology, data, workflows, workforce capability, governance, accountability and measurement must work together. These foundations are reusable – capabilities developed for one use can reduce the cost and accelerate subsequent uses. People are central to this process. Access to technology will not change how work is done unless employees have the confidence, capability and incentives to incorporate AI into redesigned workflows.

Australia has advantages in scaling AI, including strong institutions, a highly educated workforce and globally competitive industries. However, the relatively small market, concentrated industries and large number of small and medium businesses mean the path to scale will differ from that of larger economies. Large organisations can invest in the necessary foundations but often face legacy systems, regulatory requirements and operational constraints. Smaller businesses may benefit significantly from increasingly accessible AI tools, but often lack the capital, specialist skills, data and technology needed to adopt them independently. The policy challenge is therefore not only to help individual organisations, but to accelerate the diffusion of proven AI capability across the economy. This means reducing the cost and complexity of adopting proven solutions through shared infrastructure, trusted services, common standards, and reusable capabilities.

Government can accelerate the diffusion of proven AI capability by lowering the cost, complexity, and risk of adoption.

Recommendation 2: Explore how national measurement frameworks, such as the National AI Centre's AI Adoption Tracker, could better measure the depth and impact of AI adoption, including the extent to which anticipated business outcomes are realised.

Recommendation 3: Government could expand existing practical support for small businesses to access shared AI tools, industry platforms, trusted services and practical implementation support. This would allow smaller businesses to benefit from advanced AI without having to build their own technology, specialist expertise, and safeguards from scratch.

Recommendation 4: Commonwealth procurement and service delivery can help improve outcomes for Australians while establishing repeatable patterns for responsible AI deployment, measurement and assurance.

3. AI infrastructure and sovereign capability

Term of Reference (e): the opportunities for Australian sovereign AI capability, including the measures required to build, retain and commercialise domestic capability, grow research and development, and attract and retain investment and talent

Sovereign AI capability is often discussed in terms of compute, data and models. These capabilities are important, but they depend on a much larger enabling ecosystem: energy, digital connectivity, land, water, planning, capital and a workforce capable of building, operating and continually renewing AI infrastructure. Demand for AI infrastructure is already being driven globally, so the challenge for Australia is to position itself as one of the preferred locations to host that demand. Sovereign capability also extends beyond physical infrastructure to research, talent, data, intellectual property and commercial capability required to develop and retain value in Australia. Sovereign capability is also in the national interest, reducing Australia's exposure to foreign interference, disruption or withdrawal of critical AI services, and anti-competitive foreign tariffs or taxes.

Australia has a significant opportunity to establish itself as a trusted AI and digital infrastructure hub for the Asia-Pacific. However, that position is not assured. AI infrastructure capital is globally mobile and jurisdictions are increasingly competing to attract it. Australia's competitive advantages – including political and institutional stability, proximity to Asian markets, cyber and data governance, renewable energy potential, land and deep pools of infrastructure capital – will only translate into investment if projects can be delivered with sufficient speed, certainty and scale.

PwC's Global Data Centre Outlook⁵ projects that Australia could attract around \$980 billion in cumulative AI infrastructure investment between 2026 and 2050, making it the fourth most investment – intensive market in the Asia Pacific relative to GDP. Annual investment is projected to increase from \$23.1 billion in 2026 to \$58.2 billion by 2050. Capturing this investment will require timely access to power planning approvals and supporting infrastructure, while maintaining social licence by addressing community concerns about energy, water, land use, local infrastructure and local benefits.

The economic opportunity extends beyond the data centres themselves. AI infrastructure can support Australia's participation in the regional digital economy, strengthen sovereign capability and provide the compute and digital backbone for productivity growth across the wider economy.

National Cabinet's agreement in August 2026 to develop nationally consistent mandatory standards for large data centres is an important step. Standards covering energy, water, land use, infrastructure costs and skills can provide greater consistency while supporting community confidence and appropriate cost allocation. They should complement, rather than duplicate, existing state and territory planning and environmental regimes.

⁵ PwC., 2026. Global Data Centre Outlook: Where \$31.6 trillion of capex flows in the era – defining AI build – out, produced with Oxford Economics. Available at: <https://www.pwc.com/gx/en/1/services/consulting/technology/data – centre – outlook.html> (accessed 10 September 2026).

The objective should be to create a nationally coordinated framework that provides both clear obligations and clear investment pathways. Governments should seek to align energy and grid planning, land-use planning, environmental and development approvals, digital connectivity and infrastructure investment with credible developer pipelines. This would allow Australia to anticipate infrastructure requirements rather than respond project-by-project after investment proposals emerge. Our research highlights that markets which coordinate energy, planning, capital and developer pipelines are better positioned to convert underlying advantages into capacity that actually gets built. The Middle East is a contemporary example of this approach.

Sovereign AI capability also requires Australia to be able to demonstrate that AI systems developed and deployed here are safe, reliable and well governed. Investment in standards development, conformity assessment, testing infrastructure and the associated skills base can support Australian businesses seeking access to domestic and international markets. Australia should seek to contribute to the development of international practice in these areas rather than simply adopt standards developed elsewhere, and we welcome the Government's plan to develop National AI Standards.

Recommendation 5: A coordinated approach to planning, energy and investment processes would reduce regulatory complexity, improve investor confidence and help ensure Australia can deliver projects at the pace and scale required.

Recommendation 6: Recognise standards, testing and assurance capability as part of how Australia's sovereign AI capability is defined and developed.

Recommendation 7: Government to consider how Australia could build, coordinate and prioritise expertise in sovereign AI capability across research, commercialisation, standards, testing and assurance.

4. Workforce and skills

Term of Reference (g): the workforce and skills implications of AI, including the adequacy of Australia's education system, job creation and displacement, and the effects on the quality and security of jobs

Australia's workforce challenge is adapting not only to new ways of working with AI, but to AI-enabled work that was previously impossible. Jobs comprise bundles of activities that AI will affect differently. As organisations redesign workflows, the division of work between people and AI will change, with implications for roles, skills, workforce structures and job quality. Activities and workflows therefore provide a more useful basis for managing the transition than occupational forecasts alone.

Job creation and displacement: a deliberate transition

PwC's AI Jobs Barometer⁶ points to a significant reallocation of work and skills, rather than widespread net job destruction. Australian evidence to date is consistent with this.⁷ Our research found that AI job

⁶ PwC Australia., 2026. AI Jobs Barometer. Available at: <https://www.pwc.com.au/services/artificial-intelligence/ai-jobs-barometer.html> (accessed 10 September 2026).

⁷ For example, the Department of Employment and Workplace Relations' AI and Employment in Australia found no evidence of broad AI driven labour market upheaval as at February 2026: Department of Employment and Workplace Relations, 2026. AI and employment in Australia. Available at: <https://www.dewr.gov.au/workplace-relations/resources/ai-and-employment-australia> (accessed 10 September 2026).

postings in Australia doubled between 2024 and 2025, reversing four years of limited growth, while the most AI-exposed companies are growing headcount at roughly twice the rate of the least AI-exposed companies.

At the same time, AI is creating a two-track labour market. Half of Australian jobs are being ‘democratised’, as AI enables more complex tasks to be performed with lower expertise requirements, while 24% are being ‘professionalised’, as AI removes routine work and increases demand for expertise, judgement and creativity. Jobs being professionalised are growing twice as fast as those being democratised.

These changes do create transitional and equity challenges across multiple dimensions. Junior roles in which AI can assist with or perform a larger share of the work are seven times more likely to require skills traditionally associated with senior employees. Entry level roles requiring these more advanced skills grew by 35%, while entry-level roles not requiring them declined by 10%. Additionally, women are disproportionately represented in clerical and administrative occupations that are more exposed to disruption. PwC’s AI Jobs Barometer also indicates that the government and health sectors have lower AI wage premiums than some other sectors, which may warrant consideration in workforce planning as demand for AI-related skills grows. The policy task is therefore to manage the composition of the transition: protecting entry-level pathways, enabling redeployment and continuous reskilling, and ensuring that groups disproportionately exposed to disruption are not left behind.

Skills implications: dual demand for technical and human capabilities

AI is raising the value of human capabilities alongside technical ones, while skills are changing faster than organisations and the underlying training systems are adapting. PwC’s AI Jobs Barometer found that new tasks added to highly AI-exposed roles are 2.5 times more likely to rely on human – intensive skills such as judgement, empathy and creativity. AI-exposed occupations have also added an average of 187 new skills since 2019 – more than twice the rate of the least-exposed occupations. The market is not simply asking for the same work to be done faster; it is asking for different work.

This has implications for the design of Australia’s skills system. Junior positions increasingly require leadership, stakeholder management, team leadership and data-driven decision-making capabilities that education and training have not traditionally front-loaded. Technical demand is also concentrating in data analytics and interpretation, AI development and lifecycle management, cybersecurity, and human/AI interaction design. Training focused only on operating tools risks becoming obsolete before it is delivered. Skills settings should therefore build technical and human capability together and support continuous learning at the pace occupational requirements are changing.

Effects on job quality and security: emphasising human-centred AI adoption

AI can improve job quality by removing repetitive tasks, improving access to information and enabling workers to focus on higher-value activities. These benefits are not automatic: they depend on equitable access and work being redesigned to take advantage of AI’s capabilities.

The barriers are often organisational rather than technical. PwC’s AI Jobs Barometer found that only 13% of Australian businesses have performance incentives that encourage employees to experiment with and learn from AI, compared with 59% of AI leaders globally. This suggests that access without permission, incentives and work redesign is more likely to produce only incremental efficiency rather than new value.

Employers therefore have a central role to play in creating safe opportunities to experiment, recognising effective use and embedding continuous learning into day-to-day work.

We suggest the central policy objective should be human-centred AI adoption. This means protecting access to entry-level pathways, targeting capability development at cohorts and sectors most exposed to change, encouraging organisational approaches that support experimentation and continuous learning, and measuring not just adoption on its own but who benefits. Success should ultimately be assessed by whether AI improves productivity, workforce capability, job quality and economic inclusion, particularly for cohorts at risk of being left behind by technological change.

Recommendation 8: Protect entry-level pathways by encouraging education providers and employers to build judgement, communication, leadership and problem-solving capabilities earlier in workers' careers.

Recommendation 9: Support flexible, industry-informed approaches to lifelong learning and reskilling that can respond to rapidly changing occupational requirements.

Recommendation 10: Ensure public investment in AI skills develops technical proficiency alongside complementary human capabilities, including judgement, creativity, collaboration and ethical decision making.

Recommendation 11: Expand national workforce monitoring to capture participation, capability development, job quality and organisational transformation – not simply rates of AI adoption.

Recommendation 12: Incorporate AI capability into public-sector workforce planning, particularly where attracting, retaining and developing specialist skills will be critical to successful adoption.

5. Trust and responsible AI

Term of Reference (i): the adequacy of Australia's existing laws and regulatory frameworks as they apply to AI and whether there are any gaps that warrant reform.

Trust is essential to the responsible adoption of AI at scale. Clear rules, effective governance and credible assurance protect Australians while giving organisations the confidence to use AI for consequential decisions and scale what works. Trust and adoption are mutually reinforcing, and Australia can make trust a competitive advantage. Recent warnings from leading AI developers about the pace and potential severity of frontier AI risks, reinforce the need for governments to strengthen coordinated oversight, testing and assurance as capabilities evolve.

Australia's existing laws already address many of the risks associated with AI. Privacy, consumer, corporations, anti-discrimination, work health and safety, and prudential laws generally apply regardless of the technology being used. The central issue is not an absence of regulation, but that these laws were largely designed around human decision-making. Recent Australian practice shows that targeted, measures within existing frameworks can clarify how obligations apply to AI. Examples include the automated decision-making transparency obligations inserted into the *Privacy Act 1988* (Cth), and

APRA's prudential framework (including CPS 230⁸ and CPS 234⁹), APRA's April 2026 industry letter¹⁰ and joint APRA/ASIC industry roundtables.¹¹ The Voluntary AI Safety Standard has been valuable in giving organisations a common governance vocabulary, although take up will vary while it remains voluntary.

The current legal framework does not always make clear what practical compliance looks like or how an organisation can demonstrate that an AI system, or an AI agent acting for it, is meeting its legal obligations. Consistent with a technology-neutral approach, we do not consider a standalone AI Act to be the appropriate starting point. A more practical approach might be to address specific gaps in existing laws, coordinate action across regulators and clarify what organisations must do, and demonstrate, to comply. Clearer expectations and effective controls would also give organisations greater confidence to deploy AI responsibly and at scale.

The most critical gaps are:

- **AI agents and accountability.** AI agents can make transactions, grant access and act without a person approving each step. Clear rules are needed to ensure agents are identified, authorised, monitored and able to be stopped, with responsibility assigned when they act beyond their intended role. Agents could also be monitored and tested for unfair bias, including gender and cultural bias, particularly where they influence decisions affecting individuals.
- **Assurance and records.** Organisations need a consistent way to record the AI they use and show that it is operating responsibly. A recognised assurance framework, aligned with international standards, would provide boards and regulators comparable evidence while avoiding duplicate reporting.
- **Regulatory coordination.** AI can fall within the remit of several regulators, creating duplication and uncertainty. A standing coordination mechanism should set consistent economy-wide expectations for application by sector regulators. The Government's Office of AI could fulfil this function.
- **Supply-chain accountability.** Accountability generally sits with the organisation deploying AI, even where it has limited visibility of the provider or model shaping the outcome. A proportionate framework could draw on prudential regulation's concept of a material service provider, requiring organisations to identify and manage critical third-party AI dependencies while placing appropriate transparency and assurance obligations on providers.

The Committee could also consider a pre-deployment regulatory clearance pathway through which an organisation could test a proposed AI deployment with the relevant regulators before launch and receive a documented view of the expected controls and evidence. This would not remove any legal obligation, but would clarify how obligations can be met. Publishing positions in de-identified form would allow the wider market to benefit.

⁸ APRA., 2026. CPS 230 Operational Risk Management. Available at: <https://www.apra.gov.au/standards/cps-230> (accessed 10 September 2026).

⁹ APRA., 2026. CPS 234 Information Security. Available at: <https://www.apra.gov.au/standards/cps-234> (accessed 10 September 2026).

¹⁰ APRA., 2026. Letter to industry on artificial intelligence (AI), 30 April 2026. Available at: <https://www.apra.gov.au/news-and-publications/apra-letter-industry-artificial-intelligence-ai> (accessed 10 September 2026).

¹¹ <https://www.apra.gov.au/news-and-publications/insights-apra-asic-industry-roundtables>.

This model is already working overseas. The United Kingdom's pilot program AI and Digital Hub¹² enabled four regulators to provide free informal advice on questions spanning more than one remit. Its evaluation found that the cross-agency service delivered value to innovators and regulators. Utah's Office of Artificial Intelligence Policy publishes agreements made with individual participants.¹³ Other models reinforce this approach: the EU AI Act requires Member States to establish AI regulatory sandboxes,¹⁴ while Singapore's Global AI Assurance Sandbox¹⁵ supports real-world testing and assurance. Closer to home, ASIC's enhanced regulatory sandbox¹⁶ shows that Australia already has established machinery for supervised regulatory testing that could be adapted for AI.

Two design lessons stand out. First, the most valuable mechanisms coordinate several regulators because organisations' questions rarely sit within a single remit. Second, publishing outcomes allows organisations that have not applied – and incurred the associated costs – to benefit from the results. An Australian pathway should incorporate both features from the outset.

Recommendation 13: Close targeted gaps in existing frameworks first, prioritising supply chain accountability, agent identity and authorisation, and inventory and assurance obligations for high-risk uses.

Recommendation 14: Establish a standing regulator coordination mechanism (which could sit within the new Office of AI) so economy-wide AI expectations are set once and applied consistently by sector regulators.

Recommendation 15: Anchor any mandatory guardrails in a recognised assurance framework aligned to international standards ISO/IEC 42001 and NIST Risk Management Framework, and recognise continuous automated control testing as acceptable evidence.

Recommendation 16: Establish a coordinated pre-deployment regulatory clearance pathway giving organisations a documented view from the relevant regulators before launch, with determinations published in de-identified form, drawing on various international models.

Together, these measures would create a clearer pathway from innovation to trusted production. Organisations would know where accountability sits, what evidence responsible operation requires and where to obtain coordinated guidance. Over time, reusable assurance and published regulatory positions would reduce the cost of responsible adoption across the economy.

¹² Digital Regulation Co – operation Forum., 2025. Laying the foundations for cross – regulatory innovation support: insights from the AI & Digital Hub. Available at: <https://www.drcf.org.uk/siteassets/drcf/pdf – files/laying – the – foundations – for – cross – regulatory – innovation – support – – – insights – from – the – ai – – digital – hub – .pdf> (accessed 10 September 2026).

¹³ Utah Department of Commerce, Office of Artificial Intelligence Policy., 2026. AI Regulatory Relief Process. Available at: <https://commerce.utah.gov/ai/regulatory – relief/regulatory – relief – process/> (accessed 10 September 2026).

¹⁴ Articles 57 and 58.

¹⁵ AI Verify Foundation and Infocomm Media Development Authority (Singapore)., 2025. Global AI Assurance Sandbox. Available at: <https://aiverifyfoundation.sg/ai – assurance/> (accessed 10 September 2026).

¹⁶ ASIC., 2026. Enhanced regulatory sandbox (ERS). Available at: <https://www.asic.gov.au/for – business – and – companies/innovation – hub/enhanced – regulatory – sandbox – ers> (accessed 10 September 2026).