TECH POLICY
DESIGN INSTITUTE

FROM AI SOVEREIGNTY TO AI AGENCY

Measuring Capability, Agency & Power

A Practical Tool for Policymakers

DISCUSSION PAPER

NOVEMBER 2025

Level 7
221 London Circuit
Canberra ACT 2601

hello@techpolicy.au

techpolicydesigninstitute

www.techpolicy.au

ABOUT THE TECH POLICY DESIGN INSTITUTE (TPDI)

TPDi is an independent, non-partisan think tank committed to advancing best practice technology policy in Australia and globally. Based in Canberra, TPDi is registered as a not-for-profit with the Australian Charities and Not-for-Profit Commission. TPDi collaborates with all stakeholders in the tech ecosystem. Our mission is to shape technology for the benefit of humanity through rigorous research, innovative education, evidence-based public commentary, and community building.

Acknowledgments

We acknowledge the Ngunnawal and Ngambri people who are the Traditional Owners of the land upon which this report was prepared in Canberra, Australia. We pay our respects to their elders past and present. In writing about the topic of 'AI sovereignty', the authors affirm that sovereignty was never ceded.

We are grateful for the invaluable feedback we received from our peer reviewers: Andrew Brodie (Deadly Coders), Belinda Dennett (AirTrunk), Brendan Hopper (Commonwealth Bank of Australia), Clare Beaton-Wells (United Nations Youth Australia), Dave Lemphers (Maincode), David Masters (Atlassian), Elanor Huntington (CSIRO), Ian Opperman (ServiceGen), Josh Griggs (Australian Computer Society), Kate Conroy (Queensland University of Technology), Lee Hicken (National AI Centre), Liam Carroll, Alistair Reid & Tiberio Caetano (Gradient Institute), Mark Stickells (Pawsey Supercomputing Research Centre), Rosie Hicks (Australian Research Data Commons), Sally Ann Williams (Australian Academy of Technological Sciences & Engineering), Scott Winch (Sax Institute), Nadia Court, Tanya Saad & Anna Gurevich (Semiconductor Sector Service Bureau), Dr Tobias Feakin (Protostar Strategy), Simon Kriss (Sovereign AI Australia), Simon Spencer (Trideca), Dr Sue Keay (Robotics Australia Group), Taylor Hawkins (Foundations for Tomorrow) & Tim Carton (CDC Data Centres).

This project was made possible by the generous support of the Australian Computer Society and the Department of Industry, Science and Resources. All outputs represent the independent views of TPDi and/or the authors.

Generative AI tools were used in a limited capacity to assist with research and editing processes; all arguments, analysis, findings and final content were authored by the TPDi research team.

Independence statement

TPDi's independence is our most valuable asset. As a registered not-for-profit, our work is supported by external funding. We only accept funding from entities that agree to be disclosed publicly and commit to respect and promote TPDi's independence. TPDi does not represent the views of any of our funders; all outputs represent solely the views of TPDi and/or the authors.

Authors

Zoe Jay Hawkins, Co-Founder & Deputy Executive Director, TPDi & Research Associate, Oxford Internet Institute
Rebecca Razavi, Visiting Policy Fellow, Oxford Internet Institute
Meredith Hodgman, Head of Strategy & Engagement, TPDi
Johanna Weaver, Co-Founder & Executive Director, TPDi
Vili Lehdonvirta, Professor of Technology Policy, Aalto University & Senior Fellow, Oxford Internet Institute
Mercedes Page, Visiting Scholar, Massachusetts Institute of Technology Sloan School of Management & Murray Lab

This project would not have been possible without the incredible contributions of Dorina Wittmann and Heidi Brockway to the design and execution of TPDi's national consultation roadshow.

Citation

Hawkins, Z. J., Razavi, R., Hodgman, M., Weaver, J., Lehdonvirta, V., & Page, M. (2025). *From AI Sovereignty to AI Agency - Measuring Capability, Agency & Power: A Practical tool for Policymakers*. Tech Policy Design Institute. http://www.techpolicy.au/ai-agency

Contact

Tech Policy Design Institute
W: www.techpolicy.au
E: hello@techpolicy.au
ABN: 94 681 576 394

TPDI is registered as a not-for-profit with the Australian Charities and Not-for-Profit Commission.

Copyright

CONTENTS

FOREWORD

Everyone, it seems, wants 'AI sovereignty'.

Yet beneath that ambition lies confusion. The term dominates policy discussions and drives investment decisions but is used to mean everything from strategic self-reliance and resilience to cultural preservation and individual autonomy.

Behind these debates, there are more practical questions. What AI capabilities are we talking about? How do you measure them? What does sovereignty really mean in each case? Where are we today and what are the future opportunities? Only by answering these questions can we build meaningful strategies to shape AI for the benefit of humanity.

This paper proposes a shift. Rather than asking whether a nation possesses AI sovereignty, the question should be whether it has the agency to steer outcomes, protect and promote national interests, and capture value in a globally connected system.

To support this shift, the Tech Policy Design Institute (TPDi) has developed the draft **AI Agency Tool**. Informed by a **national consultation process with more than 250 experts** across government, industry, research and civil society, the draft Tool breaks down AI capabilities into clear components and provides a structured method to assess a nation's AI capabilities, agency, power and opportunity.

The Tool has been applied to the Australian context in November 2025, producing an initial snapshot of national capability and a basis for exploring future policy choices. Importantly, the Tool is designed to be adaptable, scalable and reusable, enabling its future application to other jurisdictions around the world.

The AI Agency Tool and these initial findings from its application to Australia are a work in progress and an open invitation for collaboration. Australia's AI future should be shaped by many diverse voices, not a few. We invite you to test, refine, and expand this work so that Australia can exercise true AI agency in the years ahead.

Your feedback by 15 December 2025 will inform the final iteration of the Tool, to be released in early 2026. Details on how to provide feedback and participate in this process are included in the 'Have Your Say' section of the report.

Join us to define and measure 'AI agency', so we can proactively shape a technology that is already shaping our world.

Johanna Weaver

Johanna Weaver
Co-Founder
Tech Policy Design Institute

Zoe Jay Hawkins

Zoe Jay Hawkins
Co-Founder
Tech Policy Design Institute

HAVE YOUR SAY

Australia's approach to AI should be shaped by many voices, not a few.

The draft Tool is an ambitious attempt to define the full breadth of the AI ecosystem and reframe AI sovereignty to AI agency. Its purpose is to create a shared language for policy and investment decisions, and to provide a practical resource to measure capability that can be used iteratively over time. Because of its scope, we are seeking broad and diverse participation in this next phase of refinement.

TPDi is committed to rigorous, inclusive, and transparent research. We will continue to seek to elevate historically underrepresented voices, including civil society, First Nations peoples, young people, those living with a disability, and those for whom English is a second language. These perspectives are critical to ensuring the Framework reflects the lived realities, aspirations, and values of all Australians.

How to contribute

You can provide feedback via http://www.techpolicy.au/ai-agency by **10am (Australian Eastern Daylight Time) on Monday 15 December**, on the following questions:

1. Does the Typology accurately capture the full range of capabilities, or are any missing or mischaracterised?
2. Are there any other key existing studies or assessments that could strengthen the Stocktake with further evidence of Australia's maturity in specific capabilities?
3. Does the maturity assessment accurately reflect Australia's current capability levels, or are any rated higher or lower than you would expect?
4. Does the Agency Spectrum accurately represent Australia's degree of national agency in each capability, or are any areas over- or under-stated?
5. Does the Agency Spectrum accurately capture the balance between access to international capability, domestic control, resilience through choice and export leverage?
6. Are there specific use cases or capability areas where Australia's level of AI agency should be higher than in others?
7. Are any capabilities more or less globally scarce than currently suggested?
8. Is increasing Australia's capability in a particular area more or less feasible or desirable than indicated?
9. Do the definitions and scoring systems for AI Power and Opportunity capture these concepts accurately?
10. In what ways could the Tool be enhanced to ensure it is inclusive, accessible, or actionable for your community or sector?

Next steps

Your feedback will directly inform the finalisation of the Tool and stocktake insights, which will be released in a final report in early 2026.

PART 1: DECISION MAKERS' BRIEF

The shortcomings of AI sovereignty

In a fragmenting world shaped by the rise of AI as a general-purpose technology, 'AI sovereignty' has become a central ambition of policymaking.[1] Governments are investing billions to secure it,[2] technology firms invoke it to promote expansion,[3]and civil society advocates see it as a means to safeguard rights and values.[4]

Yet the term remains contested and poorly defined.[5] This conceptual ambiguity dilutes strategy, fragments policy, blurs priorities, and leads to unproductive comparisons. In Australia, where sovereignty was never ceded by First Nations peoples, the term also carries additional significance (discussed further on page 20). Similarly, the catch-all term of 'artificial intelligence'[6] (AI) obscures the mix of capabilities that make it work, from compute and data infrastructure to AI models, skills, and governance frameworks that must be understood together for effective policy.[7]

TPDi set out to bring clarity to the debate, to define what 'AI sovereignty' means in practice, explore how it can be measured, and consider what genuine capability looks like for Australia. Details of the research method, including consultation with more than 250 experts, can be found in Part 4 of the report.

The key finding of this independent, nationally consultative research process is that AI sovereignty is a limited policy lens. Countries need a structured, evidence-based method to assess national AI capability, agency and power across the ecosystem. The AI Agency Tool provides that method.

Proposal: From AI sovereignty to AI agency & power

TPDi proposes reframing 'AI sovereignty' as 'AI agency'[8]:

AI Agency is the capacity to maintain a strategic combination of **access**, **control**, **choice**, and **leverage** over the capabilities involved in the development, use and impact of AI technologies, to steer outcomes, protect and promote personal, cultural and national interests, and capture value in a globally connected system.

While sovereignty implies ownership and control, agency offers a more pragmatic framing. It shifts the focus from defending existing assets to strengthening the capacity to act, choose, adapt, and lead.

AI agency offers a range of ways for countries beyond the US and China to pursue AI power on their own terms. It is a more pragmatic path for policy, one that recognises that many countries cannot, and need not, lead across every AI capability. AI agency involves a strategic combination domestic capability with resilient international access. It emphasises building leverage where others depend on national strengths, instead of attempting to be self-sufficient in all areas.

This strategic balance will look different for different countries. In practice, this means:

- **Maximising options:** access to resilient and diversified supply chains;
- **Reducing dependency where it matters most**: over particular capabilities and use cases; and
- **Informed decisions**: the ability to evaluate and choose from diverse options and partnerships;
- **Building leverage**: strengthening capabilities on which others depend.

Power in a connected world flows from the ability to manage relationships, not retreat from them. Technological capability as a source of national power and strategic advantage is well established.[9] Building on this approach, we recognise that power is not absolute but relational. It is shaped by what a nation can do, how independently it can act, and how rare its abilities are globally.

AI Power is defined here as:

> **AI Power** is a nation's strategic advantage in the global AI ecosystem, based on the strength of its AI capabilities (**maturity**), the capacity to act with autonomy and choice (**agency**), and the *scarcity* of those capabilities worldwide. It shows not just what a nation can do, but how independently it can act, and the leverage it gains when others depend on its strengths.

Reframing the debate toward agency and power helps decision-makers focus on where capability matters most.

The task is to understand where capability and agency can be usefully strengthened. This requires a clear and consistent method for describing and assessing AI capability, agency and power over time.

The AI Agency Tool: A self-assessment method

TPDi has developed the AI Agency Tool (the Tool), a structured and repeatable method to assess a nation's AI capabilities, agency, power and opportunity.

Part 3 of the paper describes in detail how to use the Tool, while Table 1 outlines its primary functions and benefits.

Table 1: Uses of the AI Agency Tool

Use	Function	Outcome
A common language for AI	Defines 101 AI capabilities across six layers: **the Typology**. (see Table 2)	A shared language that brings precision and comparability to national debates
Assess maturity	**Stocktake** of current capability across all layers	A curated snapshot of national strengths, areas for development and missing information
Reframes AI Sovereignty	Maps access, control, choice, and leverage through the **Agency Spectrum**	Moves from a binary notion of sovereignty to a dynamic view of agency and strategic choice
Gauge Power	Integrates maturity, agency, and global scarcity into a single view in the Power Score	Highlights areas of national advantage
Spot AI Opportunity	Weighs feasibility and desirability of capability growth in the **Opportunity Score**	Points to the most valuable and achievable areas for investment or collaboration
Inform strategy	Connects all components into one coherent policy design method	Helps decision-makers see where to build, partner, or lead
Enable transparent policymaking	Grounds decisions in evidence and clear logic	Strengthens accountability, coherence and public trust

The Typology: 101 capabilities across 6 layers of the AI ecosystem

The AI ecosystem works as a system of six layers: three technical layers (1-3) that inform the AI stack, and three enabling layers (4-6) that encircle it. Together, they cover 101 AI capabilities. Each layer is deeply intertwined with the others and with society itself. When aligned they amplify progress, when disconnected, they slow it down.

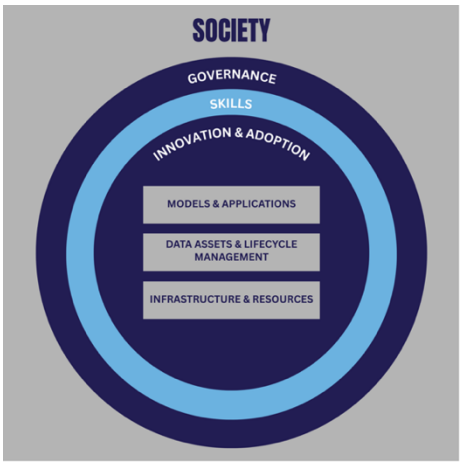

Figure 1: Layers of the AI Capability Typology

Concepts in the AI Agency Tool

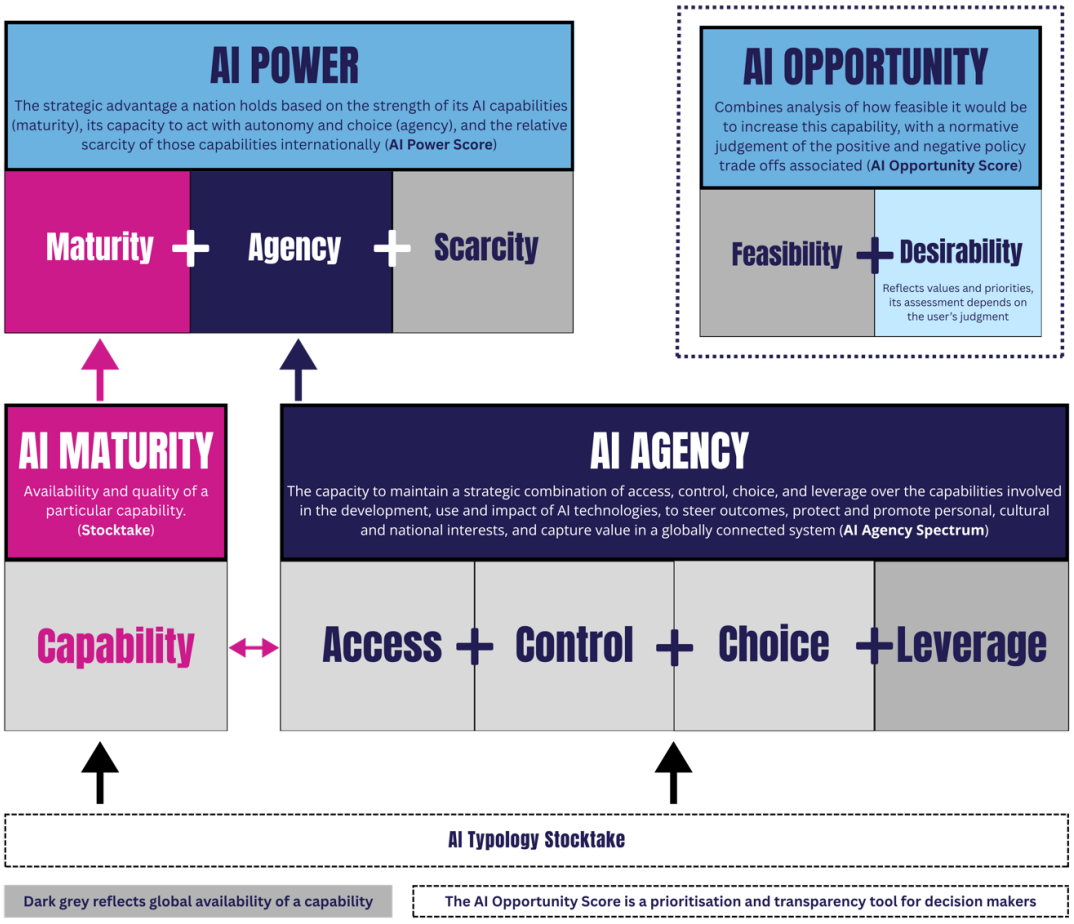

Figure 2: Concepts in the AI Agency Tool

The 'Typology at a Glance'

The 'Typology at a Glance' in Table 2 offers a structured lens through which nations can assess their maturity across 101 AI capabilities. It demonstrates how the Typology can be used to benchmark capability, identify gaps, and inform policy design.

The far-right column in Table 2 reflects the preliminary maturity findings from applying the Tool to Australia in November 2025. These findings are discussed in greater detail in Part 2 of the report.

Table 2: Typology at a Glance

AI CAPABILITY TYPOLOGY & MATURITY STOCKTAKE				
1. INFRASTRUCTURE & RESOURCES: The physical foundations of AI: compute, data centres, supply chains and natural resources.				
Category I	**Category II**	**Category III**	**Category IV**	**Australia's Maturity**
1.1 Compute & Data Infrastructure	1.1.1 Data Centres			Established
	1.1.2 Training Compute	1.1.2.1 Private Sector Training Compute	1.1.2.1.1 Cloud Training Compute Infrastructure as a Service (public cloud)	Established
			1.1.2.1.2 Private Training Compute Clusters	Emerging
		1.1.2.2 Public Sector & Public Interest Training Compute	1.1.2.2.1 Public Sector & Public Interest AI Training Infrastructure	Emerging
			1.1.2.2.2 General-purpose Public Sector & Public Interest High-Performance Compute Infrastructure	Emerging
			1.1.2.2.3 International Agreements for Cross-border Access to Training Compute	Emerging
	1.1.3 Inferencing Compute	1.1.3.1 Private Sector Inferencing Compute	1.1.3.1.1 Cloud Inferencing Compute Infrastructure as a Service (public cloud)	Established
			1.1.3.1.2 Commercial Edge Inferencing Compute Deployments	Emerging
			1.1.3.1.3 Private Inferencing Compute Deployments	Emerging
		1.1.3.2 Public Sector & Public Interest Inferencing Compute	1.1.3.2.1 Public Sector & Public Interest High-performance Inferencing Compute Clusters	Emerging
			1.1.3.2.2 Public Sector & Public Interest Edge Inferencing Compute Deployments	Established
		1.1.3.3 Consumer or Personal AI Inferencing Devices		Advanced
	1.1.4 Data Storage Infrastructure			Established
1.2 Hardware Supply chain	1.2.1 Strategic & Critical Minerals	1.2.1.1 Natural Resources		Advanced
		1.2.1.2 Extraction		Established
		1.2.1.3 Refinement & Processing		Emerging
	1.2.2 Producing Accelerators	1.2.2.1 Designing Accelerators (Fabless)		Emerging
		1.2.2.2 Manufacturing Accelerators		None
		1.2.2.3 Packaging Accelerators		None
	1.2.3 International Agreements for Accelerator Supply			Not enough data
	1.2.4 Other Critical Data Centre Hardware & Construction Inputs			Established
1.3 Supporting Infrastructure & Resources	1.3.1 Electricity	1.3.1.1 Clean Electricity Generation		Established
		1.3.1.2 Electricity Transmission & Distribution		Established
	1.3.2 Network & Connectivity	1.3.2.1 Broadband Capacity		Established
		1.3.2.2 Subsea Cables		Established
	1.3.3 Water Supply			Emerging
	1.3.4 Suitable Land			Established
	1.3.5 Permitting and Approvals Process			Established
2. DATA ASSETS & LIFECYCLE MANAGEMENT: The data required to develop and use AI: including its availability, quality of data, lifecycle, access arrangements and data sovereignty practices.				
Category I	**Category II**	**Category III**		**Australia's Maturity**
2.1 Committment to Indigenous Data Sovereignty				Emerging
2.2 Domain Specific Datasets	2.2.1 Language, Arts, Culture & History			Established
	2.2.2 Medical			Advanced
	2.2.3 Geospatial			Advanced
	2.2.4 Environment & Resources			Advanced
	2.2.5 Economic			Established
	2.2.6 Enterprise & Business Data			Not enough data
	2.2.7 Scientific, Synthetic and Simulated Research Data			Not enough data
	2.2.8 Community & Citizen Science			Not enough data
	2.2.9 Demographic			Advanced
	2.2.10 Infrastructure			Advanced
	2.2.11 Public Administration			Emerging
2.3 Data Lifecycle Management	2.3.1 Data Creation & Sourcing	2.3.1.1 Standards & Provenance		Established
		2.3.1.2 Responsible Data Sourcing		Emerging
	2.3.2 Data Preparation & Curation	2.3.2.1 Data Quality & Validation		Emerging
		2.3.2.2 Annotation & Curation (for reusability)		Emerging
	2.3.3 Data Access & Use	2.3.3.1 General Use Access		Emerging
		2.3.3.2 Availability of Government Data		Established
		2.3.3.3 Restricted Access – Copyright/IP		Emerging
		2.3.3.4 Offshore Data Access (trusted transfers)		Established
	2.3.4 Data Stewardship & Assurance	2.3.4.1 Data Retention & Archiving		Emerging
		2.3.4.2 Data Deletion & Oversight		Emerging

Table 2: Typology at a Glance

AI CAPABILITY TYPOLOGY & MATURITY STOCKTAKE			
3. MODELS & APPLICATIONS: The development and adaptation of models across technologies from computer vision to optimisation, and the applications that build on or between them			
Category I	**Category II**	**Category III**	**Australia's Maturity**
3.1 **Models**	3.1.1 **Model Development**	3.1.1.1 **Computer Vision**	Advanced
		3.1.1.2 **Computer Audition**	Emerging
		3.1.1.3 **Computer Linguistics**	Established
		3.1.1.4 **Robotics & Physical AI**	Established
		3.1.1.5 **Forecasting**	Established
		3.1.1.6 **Discovery**	Emerging
		3.1.1.7 **Planning / Optimisation**	Emerging
		3.1.1.8 **Creation / Generative**	Emerging
		3.1.1.9 **Culturally & Nationally Inclusive Models**	Emerging
	3.1.2 **Model Adaptation & Alignment**	3.1.2.1 **Domain Adaptation**	Established
		3.1.2.2 **Cultural and Linguistic Alignment**	Emerging
	3.1.3 **Model Tooling**		Emerging
	3.1.4 **Model & Agent Orchestration**		Emerging
	3.1.5 **Safety and Value Alignment**		Not enough data
3.2 **Applications**	3.2.1 **General Applications**		Advanced
	3.2.2 **Sector-specific Applications**		Established
4. INNOVATION & ADOPTION: The ecosystem of support, investment and culture that drives AI innovation, commercialisation, and adoption across society.			
Category I	**Category II**	**Category III**	**Australia's Maturity**
4.1 **Innovation**	4.1.1 **Support & Investment Availability**		Emerging
	4.1.1.2 **AI Native Companies**		Emerging
4.2 **Rate of Adoption**	4.2.1 **Private sector adoption**	4.2.1.1 **Large Enterprises**	Established
		4.2.1.2 **SMEs & Startups**	Established
	4.2.2 **Public Sector Adoption**	4.2.2.1 **Government Adoption**	Emerging
		4.2.2.2 **Defence & National Security**	Established
	4.2.3 **Public Interest Adoption**	4.2.3.1.1 **Civil Society Adoption**	Emerging
		4.2.3.1.2 **Research & Academia Adoption**	Established
	4.2.4 **Inclusive AI Adoption**		Emerging
4.3 **Culture of Adoption**	4.3.1 **Discerning Adoption**		Established
	4.3.2 **Trust in AI Deployment**	4.3.2.1 **Trust in Public Sector**	Emerging
		4.3.2.2 **Trust in Private Sector**	Emerging
		4.3.2.3 **Trust in Public Interest Sector**	Established
5. SKILLS: The skills required across all layers of the AI ecosystem, from building and developing, to governing and living with AI.			
Category I	**Category II**	**Category III**	**Australia's Maturity**
5.1 **Skills for building AI infrastructure and developing AI**	5.1.1 **Building Physical AI Infrastructure**		Advanced
	5.1.2 **Building Accelerators** (AI Chips)		Emerging
	5.1.3 **AI Research Skills**		Established
	5.1.4 **AI Development & Application Skills**		Emerging
	5.1.5 **Research and Development Capabilities** (translation)		Emerging
	5.1.6 **International AI Talent Collaborations**		Emerging
5.2 **Skills for Deploying & Maintaining AI**	5.2.1 **Business and Commercial Skills**		Established
	5.2.2 **Interdisciplinary and Domain Expertise**		Not enough data
5.3 **Skills for Governing & Securing AI**	5.3.1 **Assurance and Risk Management** (safety, bias, explainability)		Emerging
	5.3.2 **Cybersecurity and Technical Robustness**		Established
	5.3.3 **Policy and Legal Skills**		Established
5.4 **Skills for Living with AI**	5.4.1 **General Public AI Literacy and Engagement**		Emerging
6. GOVERNANCE: The Strategies, legal frameworks and policies of government and the wider ecosystem that enable trusted and influential capability.			
Category I	**Category II**		**Australia's Maturity**
6.1 **Government Strategy**	6.1.1 **National AI Strategy and Leadership**		Emerging
	6.1.2 **Policy Coherence**		Emerging
6.2 **Legal, Regulatory, Standards & Assurance Frameworks & Capabilities**	6.2.1 **Legal & Regulatory Frameworks**		Established
	6.2.2 **Ethics, Standards & Assurance Frameworks**		Established
	6.2.3 **Regulatory and Oversight Capability**		Emerging
6.3 **Institutional Capacity to Govern AI Deployment**	6.3.1 **Public Sector Institutional Capacity**		Emerging
	6.3.2 **Private Sector & Public Interest Institutional Capacity**		Emerging
6.4 **Civic Engagement and Democratic Legitimacy**			Established
6.5 **International Engagement**	6.5.1 **Influence and Norm Shaping**		Established
	6.5.2 **Access and partnerships**		Emerging

While applied here to Australia, the AI Agency Tool is intended for global use. It offers a practical framework for any country seeking to assess and benchmark national AI goals.

PART 2: AUSTRALIA TODAY

TPDi applied the Tool to produce Australia's AI Stocktake for 2025, drawing on peer-reviewed research, public data, and national consultation insights. The Stocktake provides interim findings on national strengths, dependencies and areas where more research is needed.

Interim insights for Australia

Australia's AI Stocktake Dashboard (Figure 4, page 12) summarises the findings from the Stocktake across the entire Tool. Insights on Australia's maturity by ecosystem layer are discussed on pages 13-16.

The Stocktake reveals that Australia possesses some advanced capabilities across multiple layers. However, only one layer (Infrastructure & Resources) shows predominantly established or advanced maturity; the majority of capabilities across the remaining five layers are assessed as emerging or requiring further data. Together, these layers show an ecosystem with solid foundations but uneven development and opportunities for strategic investment.

Where maturity is advanced: Australia's strengths lie in its physical and data foundations and mature data assets. It has advanced capability in the development of computer vision models and general applications. Established research, cyber security and policy and skills are combined with mature governance frameworks and international influence for a strong foundation.

Where agency is highest: Australia holds high agency over elements of the Infrastructure & Resources layer that powers AI models, from data centre development processes to critical minerals. Agency is also found to be particularly high across the Innovation & Adoption layer, which looks at the domestic uptake of AI technology across sectors and levels of public trust in AI deployment. Yet the true opportunity for each layer only emerges where agency aligns with scarcity and feasibility (as discussed on page 24).

Where maturity is emerging: Capability remains emerging in public sector and public interest AI training compute infrastructure, and other areas of compute capability. Australia demonstrates emerging model development and data lifecycle management capabilities. Research and development risk, and assurance management skills are also emerging, as is regulatory oversight. Adoption and institutional capacity to govern AI deployment across public and private sectors remains uneven.

Where agency is lowest: Australia demonstrates medium or high agency across most capabilities in the AI ecosystem, with the exception of a few including the production of accelerators and certain types of model development.

Where evidence is missing: Some capabilities could not yet be assessed due to limited data or as new categories were introduced following the consultation. Evidence gaps include international agreements for accelerator supply, proprietary enterprise and research data, and model safety and value alignment.

The Dashboard illustrates the value of breaking sovereignty down into capability, agency, power and opportunity. For example, the Innovation & Adoption layer shows that Australia has predominantly emerging maturity, but high agency. AI power only increases when agency is *combined* with a capability's maturity, which is currently uneven at this layer.

This demonstrates the importance of understanding both maturity and agency separately, but in the same frame. This is the value of the AI Agency Tool.

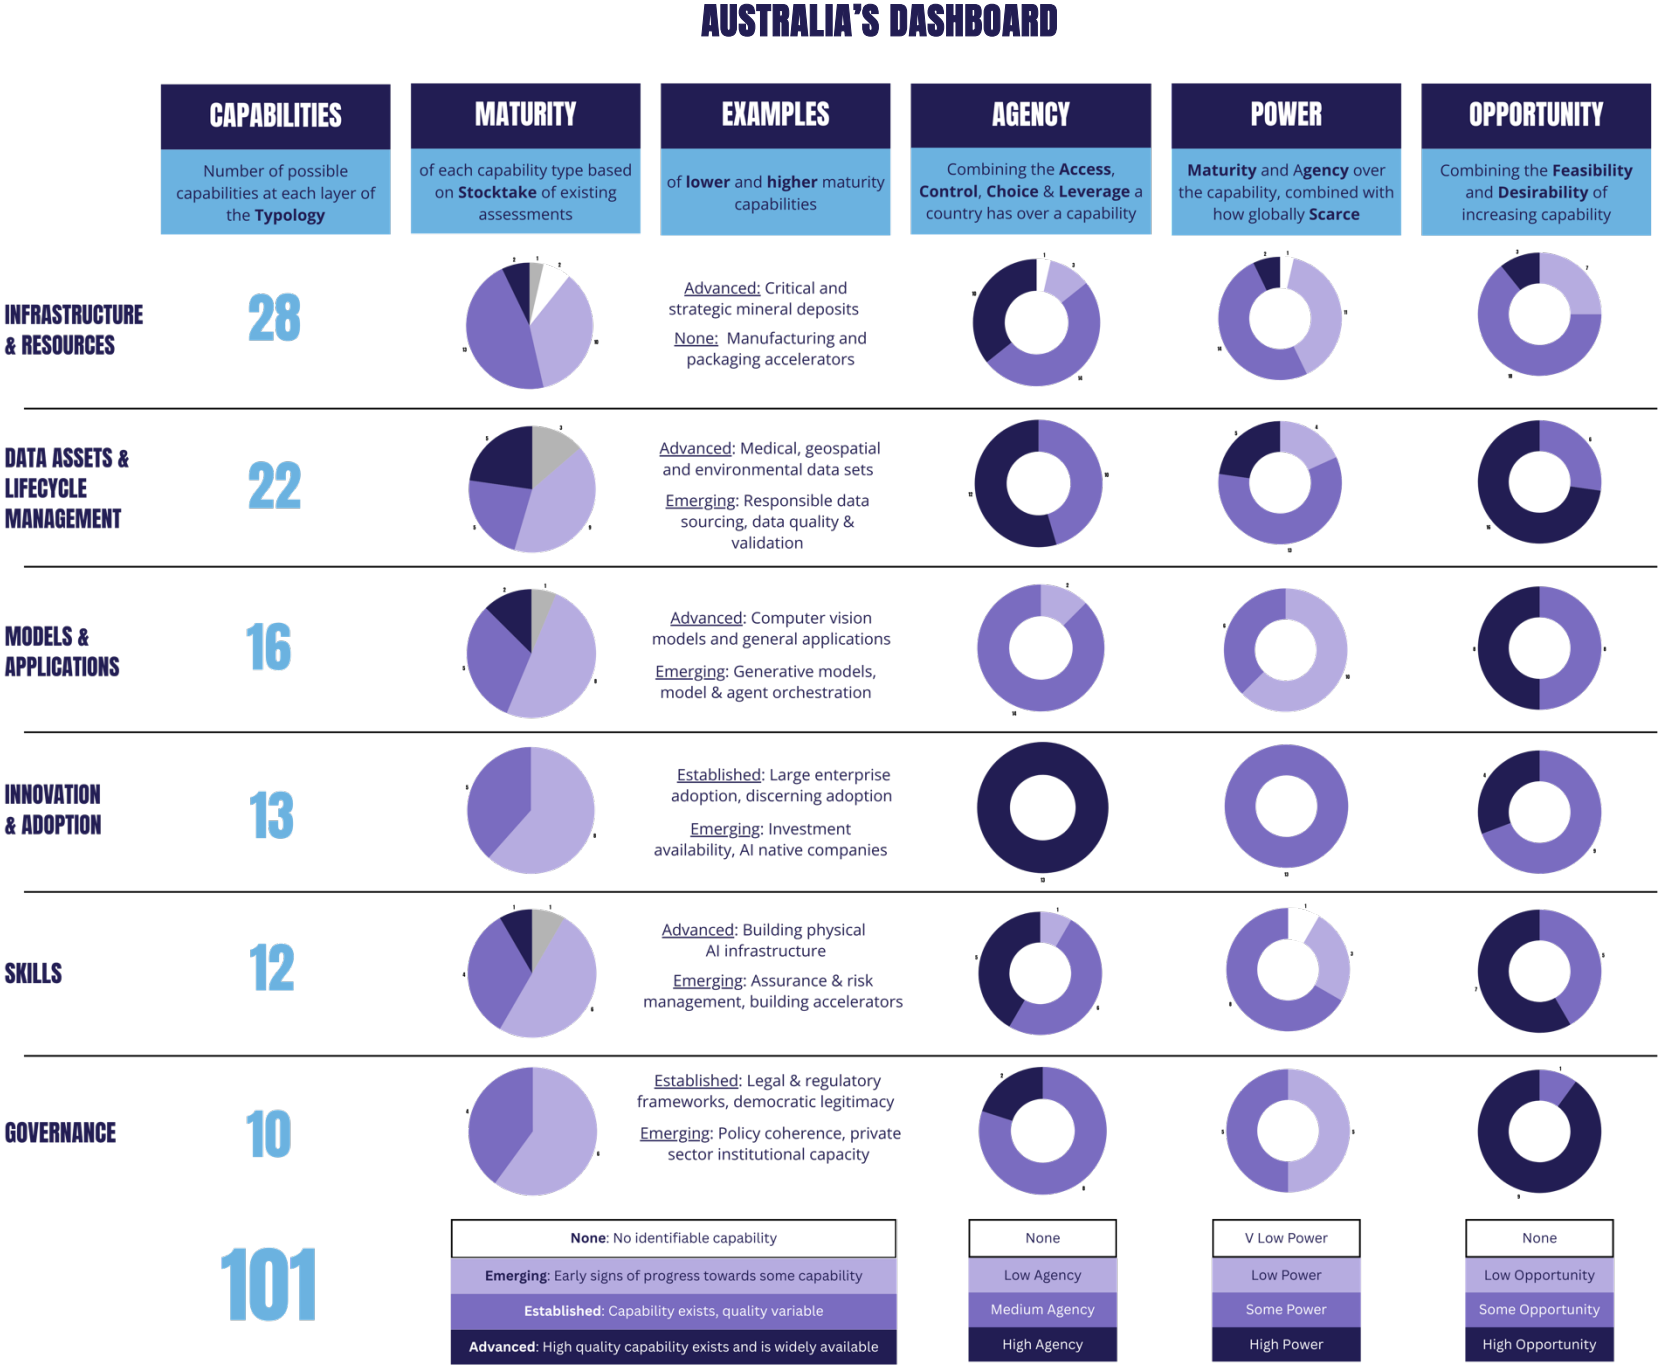

Figure 4: Australia's Stocktake Dashboard

Australian maturity insights by layer

This section draws on insights from TPDi's national AI Capability Stocktake to assess Australia's current maturity.

The following maturity assessments are best viewed in context with agency and other measures. The full Stocktake provides further details regarding our interim findings on Australia's agency, power and opportunity for each capability, and is available for download at http://www.techpolicy.au/ai-agency.

As the Tool was developed iteratively, some capability categories were not yet defined at the time of the survey or lacked sufficient evidence for assessment. The "low" or "no data" results shown here are a call to action for further input.

1. Infrastructure & Resources

AI capability begins with the physical foundations of compute power. This layer examines how effectively nations can build and sustain AI infrastructure, from data centres, training and inferencing clusters, and high-performance computing to the strategic mineral, hardware, and energy inputs that they rely on. It highlights supply chain resilience across accelerator design, production and cross-border supply, and maps the supporting infrastructure that enables compute to operate at scale, including clean energy, broadband and research networks, subsea cables, water access and suitable land. True maturity depends not only on technical strength but on whether infrastructure is established efficiently, sustainably, and in partnership with First Nations communities.

In the 2025 Australian Stocktake:

Lower maturity:

- No identifiable national capability in accelerator production and packaging.
- Emerging public sector and public interest AI training compute infrastructure for research.

Higher maturity:

- Strong energy, land, and mineral resource base underpins large-scale compute operations and sustainable expansion.
- Mature availability of consumer-grade AI powered devices and end-user inferencing capacity supports increasing AI use and adoption.

2. Data Assets & Lifecycle Management

AI systems rely on data, but data ecosystems differ widely. This layer maps the breadth, quality and diversity of national data assets across domains such as language and culture, business, health, geospatial, environment and resources, economy, and public administration. It considers how well those assets reflect a nation's reality and diversity. It highlights the importance of provenance, inclusivity and stewardship across the entire lifecycle: from creation aligned with Indigenous Data Sovereignty to ethical preparation, licensing, secure reuse and the right to delete. It also considers whether data is machine-ready, well-documented, and discoverable, and secure enough to enable responsible AI development and deployment.

In the 2025 Australian Stocktake:

Lower maturity:

- Emerging capability in responsible data sourcing as well as data quality and validation.
- Commitment to Indigenous Data Sovereignty is growing, supported by frameworks such as FAIR and CARE, though adoption remains uneven across sectors.

Higher maturity:

- Strong national datasets in health, mapping, population, and infrastructure, they're detailed, reliable, and used in research and planning.
- These are managed under clear governance and metadata systems, making them easy to share and reuse safely.

3. Models & Applications

At the heart of AI capability are the models and the applications they underpin. This layer tracks a nation's ability to develop, adapt, and deploy a range of model types, from computer vision and forecasting to robotics and generative AI. It also captures how safety, transparency and ethical alignment are embedded across model lifecycles and whether research translates effectively into real-world applications for both public and commercial use. Maturity in this layer reflects not only technical capability but the ability to align innovation with cultural and safety standards, turning ideas into impact.

In the 2025 Australian Stocktake:

Lower maturity:

- Emerging domestic development of core model architectures such as discovery, optimisation and generative AI.
- Limited evidence available to assess how safety and value-alignment are embedded across model development.

Higher maturity:

- Strong capability in computer vision model development and established capability in linguistics and forecasting model development.
- Advanced capability in the development of general-purpose applications that leverage AI for broader commercial use.

4. Innovation & Adoption

Innovation is only meaningful if it takes root and scales. This layer examines how effectively research translates into market-ready technologies, and how widely these technologies are adopted across sectors. It considers investment flows, startup activity, and pathways to commercialisation, alongside how inclusive adoption is across businesses, government, and communities, and whether the public can engage critically and confidently with AI systems. The layer reveals both the innovation engines and the social readiness that determine whether capability truly grows. Mature capability means individuals and institutions can make informed choices about whether, when and how to adopt AI, including the choice not to.

In the 2025 Australian Stocktake:

Lower maturity:

- Small but growing capability in the establishment of AI native companies, constrained by limited capital availability.
- Uneven adoption across government and civil society sectors, and adoption inclusivity across the population remains emerging.

Higher maturity:

- Established adoption across large enterprises, startups, defence, and academia anchors national diffusion capacity.
- Collaboration between research and industry is strengthening, helping new tools move from testing toward real use.

5. Skills

AI capability ultimately depends on people. This layer assesses the technical, interdisciplinary and governance skills required to design, build, deploy and oversee AI responsibly. It captures both depth and breadth, from frontier research and engineering expertise, to digital, ethical and civic literacy across the broader workforce and society. It also measures how well nations develop and retain AI talent, translate discovery between research and industry, and prepare the broader workforce to AI enabled roles. Equally important is public literacy - the ability of people, workers, and institutions to understand and engage critically with AI. True capability combines technical excellence with an informed public, creating a society able to use, question and oversee AI safety. Mature ecosystems cultivate adaptive learning systems that keep pace with technology and embed AI fluency across all sectors, including those who choose not to use it.

In the 2025 Australian Stocktake:

Lower maturity:

- Emerging assurance and risk management skills.
- Growing research-to-industry translation and international talent exchange, though still early in scale and coordination.

Higher maturity:

- Strong foundations in AI research, infrastructure engineering, cybersecurity, and policy-legal expertise.
- Business and commercial skills for AI deployment are established across key sectors.

6. Governance

Good governance, across all sectors, determines whether AI becomes a public good or a public risk. This layer examines the institutions, laws and coordination mechanisms that uphold accountability and trust. It covers national strategies, regulatory coherence, standards and assurance systems and ethical oversight alongside civic participation and international engagement. It evaluates how effectively public and private institutions govern AI with transparency and responsibility. Strong governance aligns domestic legitimacy with global influence, enabling nations to help shape the rules that shape AI. Ultimately, mature governance reflects not just compliance, but leadership in setting the terms of responsible AI at home and abroad.

In the 2025 Australian Stocktake:

Lower maturity:

- Emerging policy coherence and regulatory capacity.
- Limited private-sector and public-interest capacity for consistent AI governance, with few organisations yet adopting formal ethics, audit, or transparency frameworks.

Higher maturity:

- Legal, ethical, and assurance frameworks provide a foundation for responsible AI deployment.
- Strong civic participation and mature international engagement to shape global AI norms.

Typology in practice: Australian case studies

The AI Agency Tool is intentionally comprehensive. No single organisation or project will demonstrate every capability. The following case demonstrate how different AI capabilities in the Typology are combined and applied to specific missions, sectors, or contexts for real world outcomes. This is not an assessment or commentary on the agency of each use case.

Provided by CDC Data Centres: Project Southgate

Project Southgate is a partnership between Firmus Technologies, CDC Data Centres and NVIDIA to deliver national-scale GPU compute infrastructure powered by renewable energy. The project demonstrates activity across the Infrastructure & Resources, Innovation & Adoption, and Skills layers of the Typology, linking data centre capacity, clean energy integration, and specialist technical roles that enable large-scale compute.

Table 3: Project Southgate Case Study

Layer	Capability	In this case
5. SKILLS	5.1.1 Building Physical AI Infrastructure	Project Southgate will require advanced technical expertise to design, build, and maintain data centres and high-performance compute clusters, supporting the development of Australia's specialised AI infrastructure workforce.
4. ADOPTION	4.2.2 Public Sector Adoption	Expands domestic access to high performance GPU compute for research and public-sector experimentation within a secure domestic environment.
	4.2.1 Private Sector Adoption	Expands domestic access to high-performance GPU compute for private-sector organisations developing AI solutions in areas such as healthcare, energy, and manufacturing.
	4.1.1.2 AI Native Companies	May enable Australian AI-native firms to develop and deploy models domestically by expanding local compute access and reducing offshore dependency.
1. INFRA & RESOURCES	1.3.3 Water Supply	Utilises CDC's closed-loop LiquidCore cooling system to minimise water use in large-scale compute operations.
	1.3.1.1 Clean Electricity Generation	Underwriting of 5.1 GW in renewable power, ensuring the compute infrastructure operates on verified clean-energy sources.
	1.2.2 Producing Accelerators	Partnership with NVIDIA provides access to advanced accelerator technology and enables local integration through Firmus' AI Factory.
	1.1.3.1.1 Cloud Inference Compute Infrastructure as a Service	Project Southgate establishes national-scale AI compute infrastructure, providing local GPU access for Australian organisations to run inference and deployment workloads securely and on demand.
	1.1.2.1.1 Cloud Training Compute Infrastructure as a Service	Tasmanian founded, Firmus AI Factory platform will host model training and development workloads locally, improving energy efficiency and cost competitiveness while enabling Australian organisations to train AI systems domestically.
	1.1.1 Data Centres	CDC operates Tier 4 data centres across Australia and New Zealand, providing secure, energy-efficient, and highly reliable infrastructure for large-scale AI compute, demonstrating mature national capability in physical infrastructure.

Provided by Pawsey Supercomputing Research Centre: Setonix

Setonix, supported by the Commonwealth and Western Australian governments and operated by leading universities, is one of Australia's two Tier 1 high-performance computing facilities alongside the National Computational Infrastructure (NCI). As a leading system in the Indo-Pacific region, it demonstrates activity across the Infrastructure & Resources, Data, and Skills layers of the Typology, combining advanced compute, national research data infrastructure, and specialised scientific expertise to enable secure, large-scale AI research and collaboration.

Table 4: Setonix Case Study

Layer	Capability	In this case
6. GOVERNANCE	6.5.2 Access & Partnerships	Maintains strategic partnerships with international compute centres in Europe, the USA, and Asia to facilitate shared research capability and standards alignment.
	6.3.2 Private Sector & Public Interest Institutional Capacity	Implements responsible-AI governance practices in public research, aligning oversight and transparency with ethical and regulatory standards.
5. SKILLS	5.1.5 Research & Development Capabilities	Setonix supports large-scale AI research translating advanced modelling into scientific impact, including protein-structure and genomic analysis contributing to global datasets such as AlphaFold.
	5.1.4 AI Development & Application Skills	Pawsey runs dedicated, advanced computing training programs and uptake projects supporting code optimisation for use on advanced GPU clusters
4. ADOPTION	4.2.3.2 Research & Academia Adoption	Accelerates AI adoption across expert research communities through training, collaboration, and applied compute access.
3. MODELS & APPLICATIONS	3.2.2 Sector-Specific Applications	Supports AI-enabled research across national priority sectors including astronomy, climate science, health, and resources demonstrating cross-sector translation of AI models into applied science.
	3.1.3 Model Tooling	Provides advanced compute environments for AI-driven modelling, simulation, and large-language-model inference across research domains, customising models with specialised data.
	3.1.1.6 Discovery	Researchers identifing new patterns and hypotheses in health and life sciences, supporting disease prediction and early-stage discovery through AI-driven pattern recognition at the frontier of biomedical research.
2. DATA	2.2 Domain Specific Datasets	Provides secure data environments and management frameworks for nationally significant scientific datasets.
1. INFRA & RESOURCES	1.3.3 Water Supply	Setonix' use / resuse of aquifer water is an example of a sustainable approach
	1.3.2.1 Broadband Capacity	Supported by AARNET a NFP providing trusted high-speed, secure connectivity for national and international research institutions.
	1.3.1.1 Clean Electricity Generation	Operates among the world's most energy-efficient HPC systems, reflecting mature capability in sustainable compute infrastructure.
	1.1.4 Data Storage Infrastructure	Hosts one of the world's largest open research data storage systems, providing secure, high-capacity infrastructure for collecting and processing datasets needed for large scale AI training.
	1.1.3.2.1 Public Sector & Public Interest High-performance Inferencing Compute Clusters	Provides national high-performance compute capacity (43 petaFLOPs, 463 TB RAM) accessible to public research institutions for large-scale modelling, simulation, experimental AI inference benchmarks.

Provided by Maincode

Established to design, train, and deploy AI models entirely within Australia, Maincode is an emerging example of Australian-owned AI infrastructure development. Through its Matilda and Matilda AFL models, Maincode demonstrates shows how early-stage innovation across multiple layers of the Typology can contribute to national capability.

Table 5: Maincode Case Study

Layer	Capability	In this case
6. GOVERNANCE	6.3.2 Private Sector & Public Interest Institutional Capacity	Operates an Australian-owned AI factory platform that integrates governance, engineering, and assurance. Its internal AI Assurance Framework embeds accountability, transparency, and risk management across all model training and deployment activity.
5. SKILLS	5.1.5 Research and Development Capabilities (translation)	Bridges frontier AI research and industrial practice through an in-house team of PhD researchers and applied engineers translating new methods into production-ready Australian-made systems.
	5.1.1 Building Physical AI Infrastructure	Dedicated AI infrastructure within Australia, including MC-1 and the new MC-2 facility. Together they represent 35 million dollars of investment in industrial-grade compute and storage capacity for national AI development.
4. ADOPTION	4.2.1 Private sector adoption	Partners with Australian enterprises such as Heidi Health through pilot programs that build and test domain models on the Model Factory platform, proving how local organisations can develop Australian-made AI.
	4.2.2 Public Sector Adoption	Maincode's Model Factory provides secure, in-country infrastructure for public-interest AI initiatives, enabling government and research partners to train and operate models within Australian legal and operational control.
3. MODELS & APPLICATIONS	5.1.4 AI Development & Application Skills	Designs, trains, and deploys advanced models within Australia. The Matilda foundation model and Matilda AFL, trained on MC-1, show that large-scale AI manufacturing can be done locally to global standards.
2. DATA	2.3 Data Lifecycle Management	Maintains integrated, high-integrity data pipelines for model training and evaluation. These pipelines ensure auditability, efficiency, and compliance with Australian standards while supporting rapid model iteration.
1. INFRA & RESOURCES	1.1.4 Data Storage Infrastructure	Computing Power (Rpeak): ≈ 130 petaFLOPs Aggregate Memory: ≈ 18 TB Storage: ≈ 2.7 PB
	1.1.1 Data Centres	

Sovereignty was never ceded

Our consultation revealed another reason to shift from AI sovereignty to AI agency. The use of 'AI sovereignty' as a term in Australia warrants additional care. Here, sovereignty is not just a question of geopolitical autonomy or industrial capability, but one that intersects with an enduring and unceded sovereignty.

Sovereignty was never ceded by First Nations peoples. Framing national capability as a sovereignty issue risks obscuring the continuing sovereignties that predate Australian federation. Any national conversation about AI should reinforce, rather than distract from, the distinct and profoundly important conversations about Indigenous sovereignty and meaningfully empower Indigenous voices, leadership and agency in shaping Australia's technological future.

Insights from the dedicated First Nations consultation roundtable identified four interrelated ways to embed First Nations perspectives within Australia's approach to AI:

1. **Recognising Indigenous knowledge systems as a foundation for innovation**

First Nations peoples are this continent's first innovators.[10] Their systems of knowledge and adaptation, grounded in care for Country and Kin, reflect advanced governance, design and stewardship. These traditions emphasise balance, reciprocity and relational accountability across people, land, water, and sky. Indigenous protocols provide a robust ethical framework for technology design and deployment, embedding responsibility, transparency and respect into the architecture of innovation itself.[11] Embedding Indigenous innovation principles can strengthen Australia's pursuit of AI excellence, ensuring capability is technologically advanced, and also grounded in accountability, care, and stewardship of Country and community.

2. **Understanding whose land and resources enable AI**

AI depends on energy, minerals, land and water. Recognising these as shared resources requires meaningful consultation and consent from Traditional Owners, an ethical as well as legal foundation for sustainable capability. Mature 'capability' in these domains should include consultation and approval processes with First Nations custodians to ensure resource use aligns with environmental, cultural, and social responsibilities.

3. **Highlighting a commitment to Indigenous Data Sovereignty as a core AI tenant**

Best practice requires integrating Indigenous Data Sovereignty across all data assets and lifecycle activities. The AI Agency Tool recognises this as a distinct national capability, identifying a commitment to Indigenous Data Sovereignty through respectful handling of Indigenous Cultural and Intellectual Property and adherence to human rights frameworks and the FAIR and CARE Principles. Control of identity, knowledge, and future data use is central to sovereignty, and Australia has a short window to set these standards before AI systems become too entrenched to influence meaningfully.

4. **Continuing engagement and evolving frameworks**

The Typology adopts a Western, linear structure to support comparability. Indigenous knowledge systems are cyclical, relational and adaptive, embodying thousands of years of non-linear systems design and knowledge transmission.[12] TPDi will continue listening to First Nations Elders, technologists, and communities to evolve this work to increasingly to respect and (where appropriate) reflect Indigenous innovation traditions.

Finally, the conversation must move beyond managing risks. Building AI capability is also about creating pathways that empower First Nations technologists, entrepreneurs and communities to shape and benefit from the opportunities AI presents and to to engage confidently with AI on their own terms.

PART 3: THE AI AGENCY TOOL

Translating concepts into action: A guide to the AI Agency Tool

The draft AI Agency Tool translates the broad ideas behind AI sovereignty into a practical method for strategic decision making.

It establishes shared terms and structured assessments of national maturity, agency, power, and opportunity across 101 distinct areas of AI capability. The Tool enables government, researchers and organisations to map existing capability, identify where agency can be strengthened, and anticipate future dependencies or opportunities.

The Tool proceeds in five stages, each building on the last:

1. **Define AI capability** elements and systems with the Typology.
2. **Diagnose the current maturity** via the Stocktake.
3. **Interpret agency** using the Agency Spectrum.
4. **Analyse relative global positioning** using the Power Score.
5. **Strategise future action** through the Opportunity Score.

Together, these stages form a repeatable, evidence-based method that translates an abstract concept of AI sovereignty into clear actionable insights for decision makers.

Notes on the Tool

- **A technical evidence base**: Capability alone does not determine outcomes. Policy, market dynamics and social values all shape how technology unfolds. This Tool supports proactive socio-technical shaping. Further context on TPDi's AI policy research is available in *Tetris for Australia: Aligning our National AI Priorities*.
- **A moment in time:** the Tool captures a snapshot in time and is intended to be applied iteratively, creating benchmarks and tracking changes over time.
- **Assessments are indicative and not final:** Maturity assessments have been made based on credible, publicly available research on each capability area. These are subject to refinement as new data and insights emerge through public consultation. For transparency, the credibility of these sources refers to their methodological rigour, not their funding source. TPDi recognises the current gap in independent, publicly funded research, and notes that several reports included in the Stocktake were financed by multinational organisations.
- **Descriptive not prescriptive**: With the exception of the **Opportunity Score,** the Tool *describes* current capabilities rather than indicating what they *should* be.
- **Interdependencies matter:** Strengths create flywheels, weaknesses create bottlenecks. While the Tool doesn't map every linkage, it enables comparisons and analysis across the entire ecosystem.

The Typology: Defining national AI capabilities

To map a capability, it must first be defined. In AI policy, stakeholders often talk past one another, using the same terms to mean very different things. Without specific shared language, policymakers risk under valuing entire segments. By distinguishing between fields, such as computer vision, forecasting, optimisation, and generative AI, it becomes easier to see where strengths and emerging capabilities lie.

The AI Capability Typology brings clarity by defining and categorising 101 national elements of national AI capabilities, giving policy makers and practitioners a shared language.

The Typology organises 101 capabilities into six layers of the AI ecosystem: Infrastructure & Resources, Data Assets & Lifecycle Management, Models & Applications, Innovation & Adoption, Skills, and Governance (shown in Table 2 on page 9).

These layers are intertwined with society and work as a system. Data powers models, governance shapes adoption, and skills determine how safely AI is used. Some capabilities, such as informed citizen choice, responsible use, and public trust, span multiple layers because they link technology to social impact.

Grouped as *development capabilities* (Infrastructure, Data, Models) and *enablers* (Innovation, Skills, Governance), the Typology shows how technical foundations and social systems reinforce one another, and where national capability could grow.

The Stocktake: Measuring national maturity

To make informed policy choices, you first need to know where you stand. The Stocktake consolidates existing evidence and consultation insights to create an indicative view of maturity across Australia's national ecosystem. It represents the first attempt to bring together fragmented evaluations of the AI ecosystem into a single comparable picture.

Rather than re-assessing every capability from scratch, the Tool collates credible insights, identifies areas where evidence is thin, and highlights areas that are under-evaluated or poorly measured.

Importantly, the Stocktake is descriptive, not prescriptive. Measuring a capability's existence or maturity does not imply:

1. A value judgement: whether having more or less of a capability is inherently good or bad. These normative questions are addressed separately in the Opportunity Score, which combines the *desirability* and *feasibility* of strengthening a given area.
2. A fixed trajectory: some capabilities may plateau, evolve, or become obsolete as technologies and business models change.
3. Uniform maturity: capability levels vary widely across sectors. Additionally, some areas that appear less mature may be globally scarce, creating leverage and strategic power.

In the Tool:

- The Stocktake appears in the **pink section**.
- Supporting data can be found in each layer's corresponding 'Stocktake Sheet'.

Agency Spectrum: Reframing sovereignty

The AI Agency Spectrum translates the shift from sovereignty to AI agency into practice. It breaks down agency into four elements: access, control, choice and leverage. It recognises that power comes from balance, building domestic strength while using interdependence as a source of advantage.

The Spectrum captures the overlapping capability ownership models that co-exist within a country (international, private, public and hybrid). Rather than presenting these as distinct capability rows, the Tool layers these relationships to reveal the cumulative strategic landscape in a compact and succinct way.

Access

Access defines a country's ability to draw on international capabilities such as talent, data and infrastructure. In a globally networked AI economy, access can both be a strength and a vulnerability. Total dependence on foreign systems for critical functions risks exposure to external pressure, while strategic, diversified partnerships can create resilience.

The Agency Spectrum distinguishes between types of Access. For instance, a country may gain access to AI capabilities from jurisdictions that are governed by the *rule of law* or those where authority is exercised *extrajudicially*, such as leaders whose decisions are not subject to independent or judicial review. While both increase agency by increasing choice and resilience, access via rule-of-law jurisdictions is weighted more heavily for reliability.

For example, as it currently stands, access to international capabilities from China or Hong Kong would be represented in the 'extrajudicial reach' column, while those from the UK or US would be noted in the rule of law column. However, **none of these classifications are set and forget.** The benefit of this Tool is in its ability to be adjusted, increasing or decreasing agency in line with changes in national capability or geopolitical circumstances (for example, a rule of law country becoming subject to extrajudicial reach).

Control

Control captures the degree of domestic influence over key capabilities, whether exercised by government, research institutions, civil society or private enterprise.[13] The level of control required depends on the capability's strategic importance. For example, public interest compute, regulation, or AI safety research warrant direct domestic stewardship, while commercial applications may rely on mixed ownership models.

In this manner, having 'control' may mean that public interest research organisations possess their own AI compute training resources, that domestic AI companies are being established and grown, or that government has control of a capability, for example, regulation.

Articulating the gradient of national control is complex and central to AI sovereignty debates.

Defining what constitutes a 'local business' is often contentious. For consistency, this application of the Tool uses the Australian Government's definition of an Australian business for procurement purposes:[14]

> " a business, including any parent business, that: has 50% or more Australian ownership, or is principally traded on an Australian equities market; and is an Australian resident for tax purposes; and is a business that has its principal place of business in Australia"

Choice

Choice reflects the ability to balance access and control. A diverse and well managed mix of capability sources create flexibility, enabling fast pivots under pressure, and self-determination and adaptation as required.

In this manner, cumulative choice aids agency: the broader the options, the stronger the resilience and independence.

Leverage

If international partners rely on a country's capability for their own AI ecosystem, this creates further agency through leverage, bolstering the national negotiating position to secure or maintain access to other essential capabilities. Such dependencies amplify negotiating power and position a country as an indispensable partner in the global system.

In this case, 'leverage' involves other countries depending on Australia for the supply of critical and strategic minerals that underpin AI technologies, or commercial applications made in Australia being used in international markets, or Australia training other countries' population in particular AI related skills through our education system.

In the Tool:

- AI Agency Spectrum appears in the **grey section with the tick boxes**.
- Each layer applies these elements of Agency within their own context differently (e.g. accessing infrastructure vs. accessing skills). The logic remains constant: agency grows when access, control, choice and leverage are collectively as high as possible.

AI Power Score: Identifying competitive advantage

Power is not derived from capability alone. The **AI Power Score** integrates the maturity, agency, and global scarcity of a country's capability. It measures not only what a country can do, but how rare that ability is in an international context. This allows the Tool to highlight where a country may possess strategic leverage.

- **Maturity**: current capability levels, drawn from the Stocktake (pink column).
- **Agency**: access, control, choice, and leverage drawn from the Agency Spectrum (grey column).
- **Scarcity**: how common or rare the capability is globally, identifying potential sources of strategic advantage.

Scores are weighted to emphasise capability over scarcity, combining maturity and agency, (up to 12 points) with relative global scarcity (up to 3 points), for a total possible score of 15.

In the Tool:

- The Power Score appears in the **dark blue 'AI Power Assessment' section**.

AI Opportunity Score: Identifying options

The **AI Opportunity Score** looks forward, showing where capability could be fostered next. It assesses both the *feasibility* and *desirability* of strengthening capability, helping policy makers focus on areas that are not only achievable but worthwhile.

- **Feasibility**: how readily capability can be increased. This can shift rapidly with technological breakthrough or investment.
- **Desirability**: a holistic view of the benefits and trade-offs of expanding capability, reflecting the user's national priorities, ethics, market dynamics and public interest. Different actors will disagree, and that is the point. The score invites scrutiny and transparency by making these value judgments explicit and open to challenge. To demonstrate the functionality of this part of the Tool, this draft is populated with TPDi's subjective assessments.

In the Tool:

- The 'AI Opportunity Score' appears in the **dark blue section**.
- The Score brings together maturity, agency, power and opportunity to show how capability evolves over time.

This Tool helps decision-makers to identify where to build, where to partner, and where to lead, reinforcing the shift from a static idea of *sovereignty* to a dynamic concept of AI agency.

PART 4: METHODOLOGY

This section outlines the research process that informed the development of the draft AI Agency Tool, and situates this Discussion Paper within the broader ongoing project. It provides an overview of the project's stakeholder consultation process and demonstrates how findings from these discussions directly informed the design of the draft AI Agency Tool.

Research process

The draft AI Agency Tool was developed through an iterative, multi-stage research process combining conceptual design, expert review and national consultation.

TPDi first developed a draft framework detailing possible AI capabilities and an 'interdependence scale', drawing on desk research, including relevant TPDi research.[15] This early draft underwent targeted expert peer review to inform a revised version of the framework, which formed the basis of a national consultation.

In September 2025, TPDi undertook a comprehensive stakeholder engagement process, involving over 250 participants across Australia's AI ecosystem. The roadshow included **14 expert roundtables** across **five cities**, where experts participated in a facilitated discussion about the meaning of 'AI sovereignty' and provided direct feedback on the draft framework. Participants completed **six Stocktake Surveys**, one for each layer of the AI ecosystem, that covered the concepts behind the framework, as well as Australia's current capability levels. The themes raised in this consultation are presented on page 28.

Insights from these consultations informed the development of a draft AI Agency Tool, featuring a refined and expanded Capability Typology and new Agency Spectrum. The consultation roundtables and survey responses also served as inputs into the application of the Tool for an Australian Stocktake. This draft Tool and interim findings of the Australian Stocktake underwent further refinement following a second round of peer review on both the concepts and the Australian ecosystem findings. Peer reviewers are acknowledged on page 2.

This Discussion Paper presents both the resulting draft Tool and its preliminary findings for a second national discussion. Feedback received through this discussion process will guide the finalisation of the AI Agency Tool and Australian Stocktake, which will be released in early 2026.

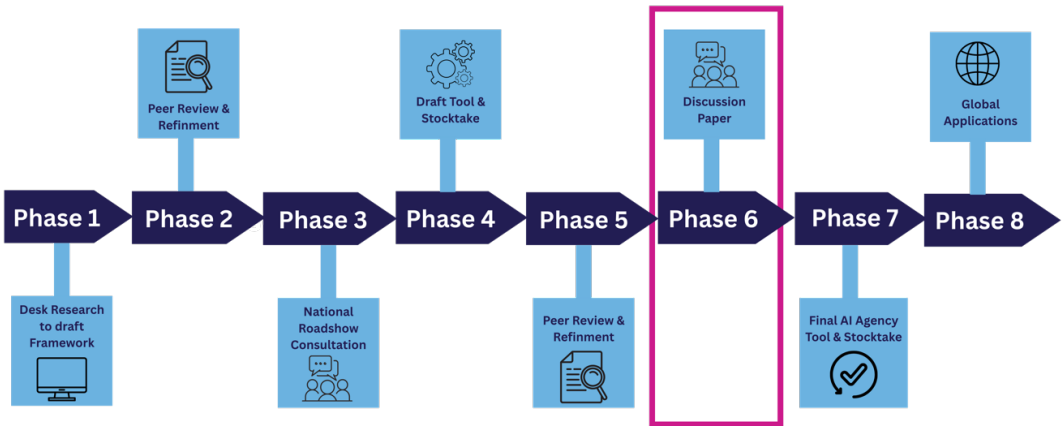

Figure 5: Timeline of research process

National roadshow stakeholder representation

The September national consultation attracted a multistakeholder cohort of participants with expertise across all elements of the AI ecosystem. Participants included representatives from civil society (7.8%), government (18%), research and education (24.2%), and industry (50%). Virtual and online surveys enabled participation beyond capital cities.

Participants represented Australia's leading digital, research, and policy institutions alongside representation from the finance, education, and creative sectors. This phase focused on expert practitioners directly involved in developing, deploying, or governing AI systems. AI users, such as SMEs and the general community were not specifically targeted in this phase.

Notably, the cohort demonstrated stronger female participation (43%) than national technology sector averages, particularly among ICT professionals.[16]

The roadshow consultation confirmed several known ecosystem gaps: limited civil-society participation (especially at the model-development layer), limited Indigenous representation despite a dedicated First Nations roundtable, and limited representation of participants under 30 years old. Ten consultation participants listed English as a second language. **Input from these systemically underrepresented groups is explicitly sought in response to this Discussion Paper.**

The activated network of national AI stakeholders represents a valuable longitudinal resource through which TPDi can sustain engagement and policy dialogue to measure progress over time.

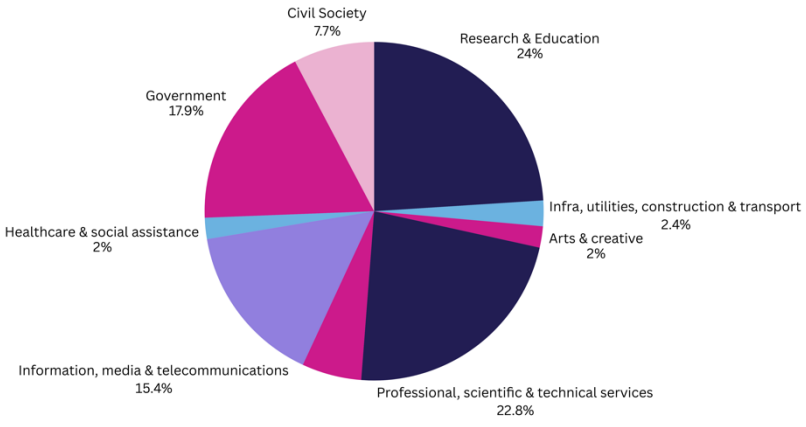

Figure 6: Pie chart of multistakeholder participation in national roadshow

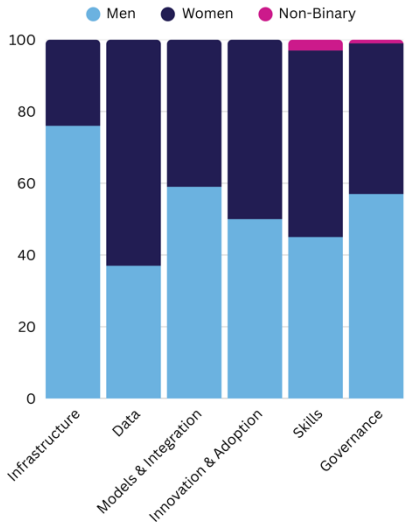

Figure 7: Graph of gender diversity of national roadshow participants across each layer

From consultation to design

Insights from TPDi's national consultation roundtables provided the foundation for the design of the AI Agency Tool.

Through structured dialogues with government, industry, civil society, and research leaders, we identified how different communities understand and pursue AI sovereignty, and why it matters in practice. These findings shaped the conceptual architecture of the Tool, translating diverse perspectives on capability and control into measurable dimensions.

Shifting from AI sovereignty to AI agency makes these distinctions visible, expanding the pathways through which a nation can build and exercise AI power, from strategic self-sufficiency and resilience to cultural alignment and the equitable distribution of benefits.

1. Autonomy & independence

Strategic self-sufficiency

Stakeholders described AI sovereignty as possessing core AI capabilities and the capacity to act independently in a volatile geopolitical environment.

> The ***Stocktake*** measures the **maturity** of 101 capabilities in the ***AI Capability Typology***. This lays the common language to describe the range of capabilities that are relevant to questions of sovereignty.

Strategic self-sufficiency in select critical capabilities, particularly national security, defence and scientific research is seen by some stakeholders as essential to reducing vulnerability to 'weaponised interdependence' (the use of economic or technological dependencies by powerful states or firms to exert pressure on others).[17]

> The ***Agency Spectrum*** reflects this logic by showing how much **control** local actors (industry, government and the public sector) have over each AI capability. It helps policymakers assess not just whether a capability exists domestically, but also the degree of control over operation, access, and governance.

Yet, total self-reliance is usually neither feasible nor desirable. The goal is balance, cultivating domestic capability while maintaining selective dependence and trusted partnerships that ensure competitiveness, innovation and access to frontier innovation.

Leverage & competitive advantage

Stakeholders identified that sovereignty is not only about reducing dependence but also about building leverage. Countries therefore strengthen their position by developing capabilities that others depend on, by cultivating competitive advantages that others may, in turn, depend on. Mutual dependence of this kind enhances strategic balance, deepens partnerships, and increases bargaining power.

This mutual dependence aligns with the concept of 'complex interdependence', which positions sovereignty as something achieved through relationships, not isolation.[18]

The ***Agency Spectrum*** maps where a country holds influence within the global AI ecosystem. It assesses not only how much control domestic actors have over specific capabilities, but also where those capabilities create **leverage** through exports and being relied upon by others. It enables policymakers to assess not just whether a capability exists domestically, but whether it can be used to counterbalance one's other dependencies.

Examples include Australia's excellence in areas such as critical minerals, renewable energy, trusted governance, and AI safety research, domains that create opportunities to export high-trust AI systems, secure infrastructure, or sustainable compute capacity. By doing so, nations reinforce their relevance in global value chains and reduce vulnerability to coercion.

Leverage is relative. The value of a country's capability increases when it is both high-impact and globally scarce.

The ***AI Power Score*** captures the maturity and agency of a country's capability, as well as how globally **scarce** such capability is. It enables policymakers to understand the national capability and agency in global context and anticipate dynamics of leverage and complex interdependence.

Participants also acknowledged how this logic of the ease or difficulty of attaining a capability plays in reverse. Whether or not a country can or should make efforts to *increase* its maturity in a particular capability is in part informed by how hard it is to do so. These constraints can be both a product of unique domestic factors and a reflection of how – and why – a capability is rare globally.

The **'Feasibility'** column in the ***AI Opportunity Score*** reflects how **difficult** it is to increase maturity in a particular capability. It enables policymakers to factor in constraints and effort required when considering where to prioritise their efforts.

2. Resilient & informed choices

Choice and evaluation of resilient systems

The capacity to sustain essential capabilities through resilience and diversification is another meaning behind conversations of AI sovereignty. Building capability through multiple sources distributes risk, reduces dependence, and strengthens flexibility. In an era of global instability, the ability to pivot quickly is a strategic asset. Sovereignty, in this sense, is not defined by ownership alone, but by options: the agency to assess, decide, and act in alignment with national priorities.

The ***Agency Spectrum*** reflects this logic by highlighting **access** to international capability as a valuable input to AI agency. It helps policymakers assess not just whether a capability exists domestically, but whether there are additional global sources of complementary, cutting-edge capability available from partners.

This echoes the sentiment in CSIRO's *Foundation Models* report which states that 'sovereign capability doesn't necessarily mean the whole AI model is developed and managed from within Australia' but rather means 'having the skills, resources and optionality to manage models built offshore.'[19]

Agency depends on having real **choice** in where AI capabilities are sourced. This dimension is captured under the ***Agency Spectrum*****,** which maps how diversified access reduces dependency and enhances national **resilience**. It helps policymakers assess not only whether a capability exists domestically, but also how multiple options can offset dependencies and strengthen national resilience.

True optionality requires evaluation literacy, the ability to judge what is trustworthy and what is not. A government can only make sovereign choices if it can understand the merits and risks of AI systems and can navigate the trade-offs they present. Building strong AI evaluation capabilities is therefore essential to exercising genuine self-determination in AI.

Individual literacy & autonomy

Participants also described sovereignty at the level of the individual. People should be able to make informed choices about how, when or whether to use AI at all. Emerging principles such as the *right to refuse* AI systems, both initially and after consent, challenge the assumption that 'more AI automatically means better outcomes'. Strengthening regulatory protections and civic literacy empowers people to question, opt out, and contest how AI shapes their daily lives, keeping human choice and autonomy at the centre of technology adoption.

Capability 4.3.1 **Discerning Adoption** within the ***Innovation & Adoption Layer***, captures the extent to which individuals can engage with AI in an informed, critical and responsible way. Greater maturity in this area enables individuals to exercise self-determination in their relationship with AI, including the right to delete.

3. Cultural alignment

AI sovereignty is not only about infrastructure, capability, or economics, but was also emphasised to be about identity, *who we are*. There is value in AI reflecting a nation's multicultural values, identities, and lived experiences, while respecting human rights, privacy and Indigenous Cultural and Intellectual Property.

Cultural data is a national asset. Datasets capturing linguistic, creative, and historical expression, from First Nations language materials and heritage archives to contemporary media and social platforms, form the foundation for models that understand and represent society authentically. When these data sources are underdeveloped or inaccessible, national identity risks being flattened into imported cultural defaults.

Growing concern surrounds the risk that AI systems trained predominantly on foreign data may dilute or distort a country's essence of voice, humour and idioms. Developing and training models domestically is as much about preserving and strengthening cultural presence in a global information ecosystem, ensuring that AI amplifies, rather than erases, the diversity of human experience.

Capability 3.1.1.9 **Culturally and Nationally Inclusive Models** and 3.1.2.2 **Cultural and Linguistic Alignment** within the ***Models & Applications Layer***, capture the extent to which models are developed and refined locally based on nationally significant data sets and local context. Greater maturity in these areas helps ensure countries do not lose cultural distinctiveness, language diversity, or national voice as AI systems become more globally integrated.

4. Domestic dividend & public good

Retain the dividend

Finally, stakeholders linked AI sovereignty to the degree to which a country can retain economic value created by AI. This includes ensuring the benefits generated by a country's data, research, and infrastructure translate into local capability, jobs and innovation instead of leaking offshore. Retaining this value means giving emerging companies an alternate path to premature foreign acquisition and preventing extractive AI models that rely on local resources while exporting profits.

To move higher in the AI value chain into model development, safety research, and applied innovation, countries need to make strategic investments in public-interest AI, transparent procurement, and innovation settings that reward domestic value creation.

Benefiting the people

AI sovereignty discussions surface a fundamental question: sovereignty for whom? Real power in the AI era is distributed, shared across governments, diverse communities, First Nations peoples, and individuals not contained within the state alone. Recognising this distribution is essential because lasting legitimacy in technology policy depends on participation and trust.

National debates often default to state-level analysis and overlook the layered and interdependent domains and groups. True sovereignty, in this view, is not centralised but shared across the society it represents. It means ensuring that AI delivers a domestic dividend that serves the public good with benefits distributed fairly across local communities rather than concentrated in government or foreign tech companies.

The '**Desirability'** column in the ***AI Opportunity Score*** measures how well increasing capability in a given area aligns with **public value** and community benefit. It helps governments, industry and civil society to assess not only whether a capability is possible, but whether it is worth pursuing.

Dialogue into design

The insights gathered through TPDi's consultation roundtables directly shaped how the AI Agency Tool was conceived and constructed. Participants' reflections on capability, interdependence, culture, and public benefit revealed that AI sovereignty is not a single goal but a dynamic balance of priorities. These perspectives informed both the structure and substance of the draft Tool, grounding abstract ideas in the lived realities, trade-offs, and aspirations expressed across Australia's AI ecosystem. In this way, the tool reflects the collective intelligence of its contributors and provides a shared foundation for the next phase of dialogue and design.

CONCLUSION

Calls for AI sovereignty increasingly shape government strategies and public debate. Yet, as this report has shown, the term is often broad, binary, and ambiguous. It tends to frame capability in terms of ownership and control, rather than in terms of the capacity to act, choose, adapt, and shape outcomes within a globally interconnected system. As a result, it offers limited guidance for practical policy design or strategic decision-making.

We propose a shift in framing from AI sovereignty to AI agency and power. This reframing supports more realistic and strategic choices. It recognises that nations do not need to lead in every capability, but require the ability to understand their strengths, reduce critical dependencies where necessary, and build leverage where national advantages exist.

The draft AI Agency Tool provides a structured and repeatable method for describing and assessing national AI capability, AI agency and AI power across the layers of the AI ecosystem. By breaking the system into clear components, the Tool supports informed decision-making and fosters a shared language across government, industry, civil society, and research communities.

The initial application of the Tool to Australia demonstrates how it can be used in practice. It shows that Australia has strong physical and data foundations, and growing technical and governance capability, but that development is uneven. Strengthening data, skills, and governance offers the most significant opportunities to increase national AI agency. Some capabilities also require further evidence, reflecting the evolving and collaborative nature of this work.

This is a work in progress. Both the Tool and the Australian findings will continue to evolve through consultation, testing, and applied use. Achieving meaningful AI agency is not the task of any single institution or sector, but a shared endeavour.

We invite you to contribute to the refinement of the AI Agency Tool and its next iteration.

By continuing to strive for greater clarity and evidence in AI policy discussions, we are better positioned to proactively shape our technological future, and the world it enables.

APPENDICES

APPENDIX 1: Defining AI

AI should be understood as an ecosystem of interlocking capabilities, not a single technology. A narrow focus on compute for generative AI overlooks the data, models, skills and governance that determine how AI is developed, deployed and controlled. Recognising and measuring these dimensions gives policymakers the breadth and precision needed to collaborate, measure progress, improve performance, and succeed.

AI extends beyond just generative models. Generative systems like ChatGPT that often dominate the headlines represent only one branch of a much larger field. Established tools such as recommendation algorithms, fraud detection systems, and automated decision-making algorithms have been transforming industries for decades. Emerging areas, including computer vision and robotics, are likely to transform society in distinct and significant ways.

For consistency and comparability, the paper adopts the Organisation of Economic Cooperation and Development's (OECD) definition of an AI system:[20]

> "An **AI system** is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment."

AI systems continue to evolve rapidly. The Typology is designed to remain relevant through these shifts, classifying model types by their functions rather than by technical methods that are likely to change more frequently.[21]

To navigate the pace of change, policy makers need a clear view of how AI capability fits together as a system. The six layers of the AI ecosystem: infrastructure and resources, data assets, models and applications, innovation and adoption, skills, and governance operate as an interdependent system. Understanding the upstream and downstream implications of each layer enables policy design to evolve alongside technology itself.

AI sits within a broader strategic technology context including quantum computing and biotech that are rapidly converging.[22] As a general-purpose technology with profound network effects, AI will both shape and accelerate these other technological advances, which have a direct impact on society. Effective policy must therefore take a joined-up socio technical view, recognising that technological capability and societal context evolve together.

APPENDIX 2: Origins of AI sovereignty

The term 'AI Sovereignty' has layered meanings and implications which complicate its use in policy debates.

'Sovereignty' is a foundational principle of the modern international system, emerging from the Treaty of Westphalia in 1648 and later formalised in international law through the United Nations Charter in 1945.[23] Although interpretations of sovereignty have changed over the years, it fundamentally refers to a state's right to govern its internal affairs without external interference from other states, establishing the principle of non-interference in international relations.[24]

However, in technology policy debates, the term 'digital sovereignty' has traditionally been used by authoritarian governments such as Russia and China to justify and advocate for state control over the internet and digital technologies.[25] In these contexts, sovereignty is used to minimise the role of non-state actors and strengthen the power of the nation-state, with serious implications for freedom of expression, privacy, and other human rights.[26] This also stands in contrast to the multistakeholder model of internet and digital governance, which emphasises the inclusion of governments, civil society, technical experts, academia and the private sector to ensure no state has complete control over the digital realm.[27]

More recently, the term 'digital sovereignty' been repurposed by many democratic governments to emphasise autonomy and reducing dependency on foreign digital infrastructure and platforms.[28] As the power and influence of big tech companies has grown, the meaning of digital sovereignty has also expanded to convey a government's efforts to counterbalance against the influence of big tech and corporations.[29] In parallel, the concept is also increasingly used to describe the autonomy and self-determination of individuals, and their ability to control their data, identities, and choices within digital systems.[30]

Within this context, 'AI sovereignty' has emerged as a subset of digital sovereignty. The term features increasingly in government policies, most recently in Canada and the United Kingdom.[31]

Definitions of 'AI sovereignty' vary. It can be broadly understood as "the capacity of a given country to understand, muster and develop AI systems, while retaining control, agency and, ultimately, self-determination over such systems."[32] NVIDIA, meanwhile, defines it through a more technical lens referring to "a nation's capabilities to produce artificial intelligence using its own infrastructure, data, workforce and business networks."[33]

With AI sovereignty gaining prominence in policy debates, there is a growing body of research examining the effectiveness, trade-offs, and unintended consequences of different policy approaches adopted in pursuit of AI sovereignty.[34]

While the term 'AI sovereignty' usefully highlights issues of control, capability, and dependency, it can imply a binary or isolationist goal that does not align with the inherently interconnected nature of the global AI ecosystem.

APPENDIX 3: Consultation list

The Tech Policy Design Institute gratefully acknowledges the contributions of the following experts and organisations who participated in workshops, surveys, and interviews informing this Discussion Paper.

Their participation reflects wide engagement across Australia's technology, research, policy and civil-society communities.

A special acknowledgment to the industry and research bodies that helped amplify the consultation process, with executive participation and opportunities for their members to contribute, including the Australian Academy of Science (AAS), Australian Computer Society (ACS), Australian Council of Learned Academies (ACOLA), Electronic Frontiers Australia Inc, IoT Alliance Australia (IoTAA), the Kingston AI Group, ARC Centre of Excellence for Automated Decision-Making and Society (ADM+S), Gradient Institute, Science and Technology Australia (STA), Tech Council of Australia (TCA), and the UNSW AI Institute.

This list extends beyond those engaged in the national roadshow and survey research, representing over 250 individuals. This broader network of 187 organisations and independent contributors provided input, feedback, or collaboration throughout the broader consultation process.

Disclaimer: Please note that participation in the consultation process does not indicate endorsement of the report's findings or recommendations.

Civil Society, Research & Academia

Adelaide University, Melissa Humphries, Jill Slay & Eddie Major
Australian Academy of Science, Lauren Sullivan
Australian Academy of Technological Sciences & Engineering, Sally-Anne Williams
Australian Council of Learned Academies (ACOLA), Prerana Mehta
Australian Institute for Machine Learning (AIML), Kathy Nicholson & Cameron Gordon
Australian National University, Sarah Logan
Australian Research Data Commons, Rosie Hicks
Brotherhood of St Laurence, Ole Chapman
Centre for Sustainable AI, Mahendra Samarawickrama
Children and Media Australia, Elizabeth Handsley
City2050, Stephen Yarwood
Curtin Institute for Data Science, Paul Hancock
Deadly Coders, Andrew Brodie
Digital Rights Watch, Lizzie O'Shea & Tom Sulston
Ethical AI Consulting, Aurelie Jacquet
Electronic Frontiers Australia, John Pane & Jonah Sullivan
Foundations for Tomorrow, Taylor Hawkins
Global Shield Australia, Devon Whittle
Good Ancestors, Greg Sadler, Emily Grundy & Luke Freeman
Good Things Australia, Jess Wilson & Linda Berrigan
Gradient Institute, Bill Simpson-Young, Alberto Chierici, Liam Carroll, Alistair Reid, Tiberio Caetano & Ali Akbari
Human Technology Institute (UTS), Nicholas Davis & Jack Goldsmith
Marconi Society, Pablo Hinojosa
Monash University, Chris Marsden
National Computation Infrastructure (ANU), Andrew Rohl

Pawsey Supercomputing Research Centre, Mark Stickells & Aditi Subramanya
Per Capita, Jordan Guiano
Queensland University of Technology, Kate Conroy, Henry Fraser & Philip Browning
RAND Australia, Austin Wyatt
Royal Melbourne Institute of Technology University, Miriam Reynoldson & Caitlin McGrane
Sax Institute, Scott Winch
Swinburne University of Technology, Elisa Bone & Linus Tan
Trust in Digital Society CRC, Maggie Lloyd
UNICEF Australia, John Livingstone
United Nations Youth Australia, Clare Beaton-Wells
United States Studies Centre, Olivia Shen & Johanna Lim
Uniting Church in Australia, Mark Zirnsak
University of Technology Sydney, Dilek Cetindamar, Jane Hogan, Asif Gill & Thom Dixon
University of Canberra, Tess Rooney
University of Melbourne, Michael Huang
University of NSW, Toby Walsh, Flora Salim, Alan Blair, Yang Song & Jiaojiao Jiang
University of Queensland, Shazia Sadiq, Michael Noetel & Alexander Saeri
University of Sydney, Kim Weatherall, Craig Jin, Rob Nicholls, Kanchana Thilakarathna & Katy Gero
University of Western Australia, Wei Liu
Victoria University, Janine Arantes & Kenea (Nea) Dhillon

Peak Bodies and Industry Associations

Australasian BIM Advisory Board, Andrew Curthoys
Australian Chamber of Commerce and Industry (ACCI), Ross Creelman
Australian Computer Society (ACS), Helen McHugh, Josh Griggs & Lynn Warneke
Australian Digital Marketing Association (ADMA), Sage Kelly
Australian Industry Group (Ai Group), David Martin
Australian Information Industry Association (AIIA), Elizabeth Whitelock
Business Council of Australia (BCA), Mike Bareja
Future Skills Organisation (FSO), Timothy Burt & Siobhan O'Sullivan
Medical Indemnity Protection Society, Lucian Burns
Robotics Australia Group, Sue Keay
Science and Technology Australia (STA), Ryan Winn
Semiconductor Sector Service Bureau (S3B), Anna Gurevich, Nadia Court & Tanya Saad
Standards Australia, Scott Brownlaw & Alex Snelson
Women in AI, Yvonne Wibowo
Working With Women Alliance, Gemma Killen

Industry

AIUC Global, Dunja Lewis
AECOM, Belinda Hodkinson
Adobe, Darshana Shah
AirTrunk, Belinda Dennett
Amazon, Mariko Lawson & Matthew O'Rourke
Amazon Web Services, Karam Kim
ARCH Cyber, Stephen Halpin
Atlassian, David Masters

Banki Haddok Fiora, Kate Haddock
Blue Crystal Solutions, Daniel Cox & Vito Rinaldi
BotWatcher, Anastasia Beasley
Brilliant Communications, Paul Wallbank
Bupa, Katelyn Nguyen
Cadent, Emma Schleiger & James Gauci
Calabash Solutions, Donna-Leigh Jackson
CDC Data Centres, Tim Carton
Commonwealth Bank of Australia, Brendan Hopper
Conor.Ai, Conor O'Neil
CSNet, Katherine Leong
CWB Consulting, Christina Wiremu-Brook
Firmus, Tim Rosenfield
Eloquium, Ishtar Vij
Empircal Legal, Nicola Cosgrove & David Turner
EY, Kim Brennan & Denise Schalet
Exchange Ecologies, Daniel Nelson
Gambarra Kaha, Michael Dickerson
Generative Boardroom, Warwick Peel
GTSystems, Rhett Sampson & Susan McHattie
H2O.ai, Trent Dolphin
Headliners Advisory, Monica Carter & Adrian J Carter
Infotech People, David Fletcher-Ly
King & Wood Mallesons, Bryony Evans
Koup Music, Kartine Ludwig
KJR, Mark Pedersen
Klapatche, Matthew Boettcher
Lucius Advisory, Kate Pounder
Macquarie Technology Group, Jamie Morse
Maincode, Dave Lemphers
Mammoth Media, David Harrison
MapAI, Adrian Towers
Mawaga Consulting, Vriti Magee
MDR Security, Rajiv Shah
Metluma, Georgie Drury
Microsoft, Kate Seward
MinterEllison, Sam Burrett, Chelsea Gordon & Jett Potter
MUZO.ai, Tim Moriarty
Montu Australia, Marta Ganko
Natirar Consulting, Peter Runcie
National Broadband Network (NBN), Ben Kereopa-Yorke
Natural Velocity, Jannat Maqbool
NextDC, Jordan Berryman
NeXtworX, Shaun Price
Nous Group, David Diviny
Old Ways, New, Angie Abdilla
Pacific Challenge Innovation Leadership, Craig Mudge
PauseAI Australia, Peter Horniak
Plumgate, Ian Christofis
Privay, Robert Postill
Project Advisory Group, Val Matthews

Protostar Strategy, Dr Tobias Feakin
PPAIA Labs, Joseph Pham
Project F, Emma Jones
Pulse Data Centre, John Henderson
QCIF, Stephen Bird
Rain Barel, Tomas Jensen
Red Marble AI, Cheryl Vize
ReadyTech, Ashvin Parameswaran
ResetData, Bass Salah & Marcel Zalloua
ROI Solutions, Karen Darling
SAP, Ryan van Leent
Scildan Legal, James Patto
Seer Data, Kristi Mansfield
ServiceGen, Ian Oppermann
ServiceNow, Simon Bush
Sharon AI, Andrew Leece & Kieran Habojan
Shape Policy Creative Intelligence, Harry Rolf
Snowflake, Rob Absalom
Sovereign Australia AI, Simon Kriss & Tiarne Hawkins
Strategic Factors, Graham Kenny
SXSW Sydney, Caroline Pegram
TechInnocens, Matthew Newman
TernaryNet, Stephen Tridgell
The Company Ethos, James Kavanagh
The Objective Company, Christian Dandre
Thoughtworks, Andy Nolan
TM Advisory, Tim Marshall
Transurban Limited, Angela McGinness
Threat Canary, Andrew Horton
Trideca, Simon Spencer
TUTXI, Jenny Wu
V & Co, Toby Vervaart
Vantage Data Centres, Quynh Do & Martin Kazimier
VAZ, Keith Vaz
Whiley Group, Angkana Whiley
Xero, Grace Gown
Xelleron Pty Ltd, Stefan Slomka

Government Departments & Agencies

Attorney-General's Department
Australian Digital Health Agency (ADHA)
Australian Research Data Commons (ARDC)
Australian Signals Directorate (ASD)
Australian Taxation Office (ATO)
Department of Climate Change, Energy, the Environment and Water (DCCEEW)
Department of Defence
Department of Finance
Department of Foreign Affairs and Trade (DFAT)
Department of Health
Department of Home Affairs

Department of Industry, Science and Resources (DISR)
Department of Infrastructure, Transport, Regional Development, Communications, Sports and the Arts
Department of the Prime Minister and Cabinet (PM&C)
Digital Transformation Agency (DTA)
Customer Service New South Wales
Commonwealth Scientific and Industrial Research Association
Legal Aid New South Wales
National Disability Insurance Agency (NDIA)
National AI Centre
Office of National Intelligence (ONI)
Queensland State Development
QLD Police Service
Services Australia
South Australian Department of State Development
Treasury
Victorian Government Department of Government Services
Whitehorse City Council, Victoria

Independent and Individual Contributors

Martin De Domenico
Angelica Chowdhury
Brian Hanlon
David Harding
Sylvie Hu
Jean Linis-Dinco
Melissa Misuraca
Will Stoltz

REFERENCES AND NOTES

[1] Australian Academy of Technological Sciences and Engineering. (22 August 2025). *Made in Australia: Our AI opportunity*. https://atse.org.au/what-we-do/strategic-advice/made-in-australia-our-ai-opportunity/.; Chavez, P. (7 November 2024). *Sovereign AI in a hybrid world: National strategies and policy responses.* Lawfare. https://www.lawfaremedia.org/article/sovereign-ai-in-a-hybrid-world--national-strategies-and-policy-responses.

[2] European Commission. (October 2025). *Apply AI Strategy.* https://digital-strategy.ec.europa.eu/en/policies/apply-ai.; Department of Science, Innovation and Technology (United Kingdom). (January 2025). *AI Opportunities Action Plan*. https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan.; Innovation, Science and Economic Development Canada (2024). *Canadian Sovereign AI Compute Strategy*. https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy.

[3] Althoff, J. (16 June 2025). *Announcing comprehensive sovereign solutions empowering European organizations*. Microsoft Blog. https://blogs.microsoft.com/blog/2025/06/16/announcing-comprehensive-sovereign-solutions-empowering-european-organizations/.; Garman, M. (November 2023). *AWS Digital sovereignty pledge: Control without compromise.* AWS Security Blog. https://aws.amazon.com/blogs/security/aws-digital-sovereignty-pledge-control-without-compromise.

[4]Gornemann, E. (2024). *Digital Sovereignty (Fundamentals series).* Berlin: Weizenbaum Institute. https://fundamentals.weizenbaum-institut.de/en/digital-sovereignty/.

[5] See Part 4 of this report (page 25) for insights into the research process, including details of the national consultation.

[6] OECD. (2024). *Explanatory memorandum on the updated OECD definition of an AI system (OECD Artificial Intelligence Papers, No. 8)*. OECD Publishing. https://doi.org/10.1787/623da898-en.

[7] See Appendix 1 of this report for discussion on defining AI.

[8] While 'AI agency' is often used to describe the autonomy of machines, here we refer to human and national capability. This is closely linked to the growing commercial market for 'AI agents'. In these conversations, agency is used to talk about the independence and influence of AI technologies themselves. See Kang, H., & Lou, C. (2022). AI agency vs. human agency: Understanding human–AI interactions on TikTok and their implications for user engagement. *Journal of Computer-Mediated Communication, 27*(5). https://doi.org/10.1093/jcmc/zmac014; Legaspi, R., He, Z., & Toyoizumi, T. (2019). Synthetic agency: Sense of agency in artificial intelligence. *Current Opinion in Behavioral Sciences 29,* 84–90. https://doi.org/10.1016/j.cobeha.2019.04.013.

[9] Waltz, K. N. (1979). *Theory of international politics.* Reading, MA: Addison-Wesley; Gilpin, R. (1981). *War and change in world politics.* Cambridge: Cambridge University Press; Strange, S. (1988). *States and markets.* London: Pinter.

[10] McNiven, I., & Russell, L. W. (2023). *First Knowledges Innovation: Knowledge and Ingenuity* (1st ed.). Thames & Hudson.

[11] Abdilla, A; Kelleher, M; Shaw, R; Yunkaporta, T. (2021). *Out of the Black Box: Indigenous Protocols for AI.* Deakin University. https://hdl.handle.net/10536/DRO/DU:30159239.

[12] Yunkaporta, T. (2019). *Sand talk: How Indigenous thinking can save the world*. Text Publishing. https://cat2.lib.unimelb.edu.au:443/record=b7338253~S30.

[13] Rosenau, J. N. (1976). Capabilities and control in an interdependent world. *International Security, 1(2),* 32-49.

[14] Department of Finance (2025). *Guidance on the definition of an Australian of New Zealand business.* www.finance.gov.au/government/procurement/defining-australian-business-commonwealth-procurement/definition-australian-business.
[15] Hawkins, Z. J., Lehdonvirta, V. & Wu, B. (20 June 2025). *AI Compute Sovereignty: Infrastructure Control Across Territories, Cloud Providers, and Accelerators*. SSRN. https://ssrn.com/abstract=5312977.
[16] Compared with most recent ABS data breakdown available: Australian Bureau of Statistics (2022). *2021 Census data shows Australia going high tech*. https://www.abs.gov.au/media-centre/media-releases/2021-census-data-shows-australia-going-high-tech.
[17] Farrell, H. and Newman, A. (2019). Weaponized Interdependence: How Global Economic Networks Shape State Coercion. *International Security44 (1),* 42–79. https://doi.org/10.1162/isec_a_00351.
[18] Keohane, R.O. and Nye, J.S. (1977). *Power and Interdependence: World Politics in Transition*. Boston: Little, Brown and Company.
[19] CSIRO. (2024). *AI foundation models: Australia's opportunity in the age of generative AI.* https://www.csiro.au/en/research/technology-space/ai/ai-foundation-models-report.
[20] OECD. (2024). *Explanatory memorandum on the updated OECD definition of an AI system (OECD Artificial Intelligence Papers, No. 8)*. OECD Publishing. https://doi.org/10.1787/623da898-en.
[21] European Institute of Innovation and Technology (2021). *Creation of a Taxonomy for the European AI Ecosystem*. https://www.eit.europa.eu/library/creation-taxonomy-european-ai-ecosystem.
[22] Suleyman, M., & Bhaskar, M. (2023). *The coming wave: technology, power, and the twenty-first century's greatest dilemma* (1st ed.). Crown; Department of Industry, Science and Resources. (19 May 2023). *List of Critical Technologies in the National Interest.* https://www.industry.gov.au/publications/list-critical-technologies-national-interest.
[23] Croxton, D. (1999). The Peace of Westphalia of 1648 and the Origins of Sovereignty. *The International History Review, 21(3)*, 569–591; United Nations. (1945). *Charter of the United Nations, Chapter I: Purposes and Principles.* https://www.un.org/en/about-us/un-charter/chapter-1.
[24] Bull, H. (1995). *The anarchical society: A study of order in world politics* (2nd ed.). Macmillan.
[25] Information Office of the State Council of the People's Republic of China. (June 2010). *The Internet in China (White Paper).* http://un.china-mission.gov.cn/eng/chinaandun/economicdevelopment/kj/201010/P020210901031412447228.pdf.; Ministry of Foreign Affairs of the People's Republic of China. (March 2017). *International Strategy of Cooperation on Cyberspace* . https://www.fmprc.gov.cn/mfa_eng/wjb_663304/zzjg_663340/jks_665232/kjlc_665236/qtwt_665250/201703/t20170301_599869.html.
[26] Sherman, J. (2022). *China's war for control of global internet governance*. SSRN. https://ssrn.com/abstract=4174453.
[27] Internet Society. (27 May 2024). *It Takes a Community to Defend the Internet* . https://www.internetsociety.org/blog/2024/05/it-takes-a-community-to-defend-the-internet/.
[28] Pohle, J., & Thiel, T. (2020). Digital sovereignty. *Internet Policy Review, 9(4).* https://doi.org/10.14763/2020.4.1532.
[29] Van Dijck J. 2020. Governing digital societies: Private platforms, public values. *Computer Law & Security Review* 36: 105377. https://doi.org/10.1016/j.clsr.2019.105377.
[30] Belli, L., & Gaspar, W. (October 2023). *The quest for AI sovereignty, transparency and accountability.* UN IGF Data and Artificial Intelligence Governance Coalition. https://www.intgovforum.org/en/filedepot_download/229/26953.
[31] Government of Canada. *Canadian Sovereign AI Compute Strategy*. https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy

[32]Belli, L., & Gaspar, W. (October 2023). *The quest for AI sovereignty, transparency and accountability.* UN IGF Data and Artificial Intelligence Governance Coalition. https://www.intgovforum.org/en/filedepot_download/229/26953.

[33] NVIDIA. (28 February 2024). *What Is Sovereign AI?* NVIDIA Blog. https://blogs.nvidia.com/blog/what-is-sovereign-ai/.

[34] Stark, J. & Gill, A. From Funding to Trust: Australia's Sovereign AI Architecture. *TechRxiv.* September 15, 2025. DOI: 10.36227/techrxiv.175790355.58552311/v1

TECH POLICY
DESIGN INSTITUTE