

Australia's AI Dividend

Shared ownership, public-private partnering
and revenue apportionment for
AI training infrastructure

AUSTRALIA'S AI DIVIDEND

Shared ownership, public-private partnering and revenue apportionment for AI training infrastructure

Authors: Michael D’Rosario, Wesa Chau, Peter Lewis
September 2026

This report has been prepared by Per Capita. It is informed by a roundtable process Per Capita conducted with a range of sector stakeholders, including research teams, industry and professional associations, sector organisations and individual experts. The report does not necessarily reflect the views of any of those participants, and no participant should be taken to endorse its analysis or recommendations.

Per Capita’s research team examined the plausible approaches to capturing the national interest in AI training infrastructure through the lens of the UK Government Office for Science’s AI 2030 scenario framework, and the recommendations set out above represent one analytically defensible position intended to be tested in good faith against alternatives.

Acknowledging the fast-changing nature of AI and aligned sectors, and associated sectoral developments, the information presented is as accurate as at the time of writing.

This report was written by a human and refined with the assistance of Claude Opus 4.8/5.

Table of Contents

<i>Table of Contents</i>	3
<i>1 Executive Summary</i>	5
1.1 Summary of recommendations	8
<i>2 Introduction</i>	10
<i>3 Designing for an uncertain future</i>	13
3.1 Examining the common factors	15
3.2 A VUCA challenge calls for a balanced response	16
<i>4 The opportunity and the window</i>	16
4.1 The policy window remains open, but leverage will shift	17
4.2 Why the window is closing	18
<i>5 What other jurisdictions are doing</i>	19
<i>6 Why a single instrument fails</i>	21
<i>7 The recommended model: layered capture through partnering and shared ownership</i>	23
7.1 Layer one: shared ownership within a partnership framework	23
7.2 Layer two: revenue apportionment anchored to countable units	24
7.3 Layer three: an access fee recognising the public endowment	26
7.4 Layer four: reserved sovereign compute	26
7.5 Layer five: a training and apprenticeship contribution	27
<i>8 Stress-testing the model against the 2030 scenarios</i>	27
<i>9 Governing the proceeds</i>	30
<i>10 Protecting the creative sector</i>	31
<i>11 Conditions on the training privilege</i>	32
11.1 A commitment to local talent and apprentices	33
<i>12 Trade-offs and the case for ambition</i>	34
<i>13 Conclusion</i>	36
<i>Appendix A: The Attributable Influence Principle</i>	37
References (Appendix A)	40


<i>Appendix B: Sustaining human knowledge production in the AI economy</i>	41
B.1 The circular flow of the knowledge economy	42
B.2 Information quality and market failure	42
References (Appendix B)	43
Notes	43

1 Executive Summary

Investment in AI infrastructure continues to rise rapidly, with technology companies committing more than US\$650 billion to AI infrastructure in 2026.¹ Within the U.S. it represents the second largest annual investment in history, since the railroad network. The scale of the interest in investment is represented in Figure 1. While this interest is both cause for consternation and optimism it is apparent that this is a significant moment. It is anticipated (and the current construction pipelines suggest) that the relative level of investment in Australia will be (broadly) comparable, albeit smaller (as a percentage of GDP). The scale of investment is significant greater than any recent investment in domestic infrastructure notwithstanding the mining boom.

Notably interest exists in expanding Australia's inference asset base but also in the strategically important compute infrastructure base. AI training infrastructure, often called 'compute' infrastructure, is much more concentrated in a small number of locations than the computing infrastructure used to run everyday AI applications. This is because building and operating this infrastructure requires significant investment, energy, and resources, many countries may not be able to host it at scale.

Figure 1 – Selected investment 'boom' intervals




Note: The y axis denotes the scale of investment to the relative trough. Source: Bank of International Settlements 2026

¹ Matt Day, 'How Much Is Big Tech Spending on AI Computing? A Staggering \$650 Billion in 2026', *Bloomberg*, 6 February 2026.

Australia is well positioned to become a major hub for AI compute infrastructure. It offers several advantages, including abundant renewable energy resources, political and regulatory stability, a highly skilled workforce, and a strategic location close to Southeast and South Asian markets.

Figure 2 – Australian infrastructure investment (selected) as a percentage of GDP



Note: This analysis represents Per Capita's estimate a plausible central estimate of investment. Source: Productivity Commission data, ABS datasets, Per Capita Analysis


These unique advantages mean Australia currently holds important bargaining power, however, they are only temporary. Once infrastructure, grid connections, and supply chains are in place, leverage typically shifts to the firms that own and operate the assets, as seen in both international experience and Australia's offshore gas industry.

This represents a significant opportunity for Australia, but also a source of uncertainty, given the variety of possible approaches and the range of potential technological trajectories. Given the fast-changing nature of the developments of AI, therefore any policy framework for AI training infrastructure must be designed to operate under conditions of substantial volatility, uncertainty, complexity and ambiguity of the sector. There are five scenarios encompass a wide range of possible outcomes, with correspondingly different implications for policy design:

1. 'Slow Burn': progress slows after 2028, adoption remains constrained and AI produces relatively modest economic gains, with labour displacement concentrated in a limited number of occupations.
2. 'Open Frontier': technological progress also slows, but open-weight systems narrow the gap with frontier proprietary models, distributing both productive opportunities and the capacity for harmful use more widely.

3. 'Augmented Growth': capability development continues while humans remain substantially 'in the loop', productivity increases, new forms of employment offset displacement, and the resulting economic gains are relatively broadly distributed.
4. 'Transformation Economy': produces a markedly different distributional outcome, comparable technological progress occurs, but human labour is displaced more extensively and a substantial share of the resulting profits accrues to a small number of foreign frontier firms.
5. 'Take-Off': AI systems surpass expert human performance across almost all cognitive tasks within a comparatively short period, economic growth accelerates substantially, and an increasing share of value accrues to the owners of compute, intellectual property and capital within an increasingly 'K-shaped' economy.

Figure 3 – The economic impact of AI under different uptake and change scenarios



Note: This is a conceptual diagram presenting the plausible impact of AI in society. Derived from UK GO Science scenarios 2026.

The variable impacts are most easily understood when considering Figure 3. Therefore, recognising the wide range of views on the direction, pace, and economic impacts of its development, and acknowledging the volatility, uncertainty, complexity, and ambiguity that characterise it, the following recommendations have been developed to ensure everyday

Australians benefit from the investments and development at a time Australia retains negotiating leverage.

The key question is not whether AI infrastructure should operate in Australia, but on what terms of its operation, so that the economic value generated flows back to Australians.

Per Capita's contribution to this discussion to highlight approaches that maximising the distribution of the potential economic benefits in these scenarios, while mitigating the concentration of the economic impacts.

Working with civil society stakeholders we conclude that rather than a single 'silver bullet' economic policy such as a specific taxation instrument or redistributive tool, the government needs a sequenced and integrated approach to asserting our economic sovereignty.

1.1 Summary of recommendations

1. Adopt a **layered capture model** rather than a single instrument. No one mechanism is both ambitious and robust; the strengths of each offset the weaknesses of the others. Layered capture will ensure that a balanced dividend is received, without the risk of a forecast scenario not coming to fruition.
2. Choose the model for **robustness across futures**, not for any single forecast. Because AI's trajectory is deeply uncertain, the test of a good design is that it captures value whether progress stalls or accelerates, not that it is optimised for the future we happen to expect.
3. Take a **public equity stake** in the fixed enabling assets (data centre campuses, compute facilities, transmission and the renewable generation behind them) through existing vehicles such as the Clean Energy Finance Corporation, and NBN Co, or a new entity, negotiated before commitment, when stakes are cheapest. Emphasis must be placed on the balance of assets, not staking all investment just in compute.
4. Structure investment **as genuine public-private partnerships** with public co-investment, board representation, transparency and step-in rights, not as passive grants or tax concessions. This approach ensures that Australia is a genuine participant and rather than a bystander in the process.
5. Apply an **ongoing revenue apportionment** anchored to countable licence units (registered models, deployment instances) rather than to accounting profit, so that the charge scales with commercial success but cannot be eroded by transfer pricing in the way the Petroleum Resource Rent Tax was. The countable unit is a present-day proxy for the contribution that Australian inputs

make to a model's value, the principle of attributable influence developed in Appendix A, to which the charge could in time be more directly calibrated.

6. Reserve a **share of sovereign compute** in kind for universities, public-interest research, start-ups, the national science agency and strategic needs, as a condition of access to Australian infrastructure.
7. Quarantine the proceeds in a **dedicated national fund** on the fiscal architecture of the Future Fund or Ireland's Future Ireland Fund, separating long-term asset value from the operating budget.
8. Build a **creative-sector mechanism** for structural displacement that is distinct from copyright licensing, recognising that copyright, including the strengthened creator protections the Government has flagged, addresses the licensing of works, while the structural erosion of creative livelihoods needs a separate response; and attach enforceable national-interest conditions to the training privilege.
9. Build on the legislated skills baseline with a genuine commitment to **local talent and apprentices**. With training and apprenticeship requirements set to become part of the mandatory data centre standards, the question is how to make them meaningful; an apprentice-to-tradesperson ratio attached to grid-connection and planning approvals is one option worth serious consideration, with the precise settings developed in dialogue with training entities, the TAFE sector, group training organisations and unions, and matched by public investment in training capacity.

2 Introduction

Investment in AI infrastructure continues to rise rapidly, with technology companies committing more than US\$650 billion to AI infrastructure in 2026.² Within the U.S. it represents the second largest annual investment in history, since the railroad network. It is anticipated (and the current construction pipelines suggest) that the relative level of investment in Australia will be comparable (as a percentage of GDP). This presents as a significant opportunity, but also as a genuine source for consternation, given the many plausible approaches that could be taken to the opportunity, and many trajectories for the technology. In March 2026 the New South Wales Investment Delivery Authority endorsed fifteen data centre projects worth around A\$52 billion for prioritised approvals support, and the Department of Industry has noted that announced investments could scale beyond A\$100 billion nationally.³

AI training infrastructure (compute infrastructure), the infrastructure used to develop frontier models, is far more geographically concentrated than the inference computing infrastructure that runs everyday applications. It is plausible that many countries will not be able to host the necessary infrastructure, given the mix of investment and resources required. Australia possesses unique geographic and resource advantages, making it a highly favourable location for 'compute' infrastructure, abundant renewable energy (assuming the plausible leveraging of this advantage, in this moment), stable governance and a sound legislative framework, a highly educated workforce, and geographic positioning proximate to South East (and South) Asian markets. Notably Australia also maintains a strong copyright framework with no text and data mining exception. That last point matters greatly, presenting as both a challenge and an opportunity. In practical terms it means AI companies cannot lawfully train frontier models on protected material without permission, so Australia retains genuine leverage over the terms on which training happens here.

This report does not propose changes to Australia's copyright provisions and should not be understood as a proposal for copyright compensation. It does not seek to alter the mechanisms through which creators are remunerated for the use of their work. Australia already has an established copyright framework governing the licensing of creative works, and the Government has signalled its intention to strengthen protections available to creators where their work is used in AI training, including greater control over the terms and value of that use. Moreover, in the absence of a text and data mining exception, lawful use of protected material for AI training remains subject to the permissions required under the existing copyright framework.

² Matt Day, 'How Much Is Big Tech Spending on AI Computing? A Staggering \$650 Billion in 2026', *Bloomberg*, 6 February 2026.

³ NSW Government, 'NSW Government Backs Data Centre Investment, Sets Course for Sustainable Development', 27 March 2026.

The subject of this report is separate and additional. The public and national interest in the economic value created when frontier AI systems are trained in Australia using Australian energy, land, infrastructure, public investment and institutional stability. Capturing a share of this 'national dividend' is principally a question of industrial and fiscal policy rather than copyright remedy, and the two issues should not be conflated. Nor is it necessary to diminish the interests of copyright holders in order to secure a broader public dividend. A policy framework can recognise the legitimate interests of rights holders, technology developers and the wider community simultaneously. Where this report considers the creative sector, it therefore does so beyond the core function of copyright. In particular, it addresses the potential for structural economic displacement within creative industries, a problem that copyright licensing was not designed to remedy, and which is analytically distinct from compensation for the use of protected works.

The report recommends that Australia treat AI training infrastructure not simply as private investment to be facilitated or regulated, but as strategically significant infrastructure in which the public should hold a durable and meaningful participating stake. Its central argument is that no single policy instrument, whether a tax, access fee or equity holding, is likely to secure the national interest across the range of plausible technological and economic outcomes. Each instrument has characteristic strengths and limitations, and each performs differently under alternative future conditions. The recommended response is therefore a layered model combining public-private partnership and shared ownership of fixed enabling assets with ongoing revenue apportionment linked to the commercial value those assets help to create. These fiscal and ownership mechanisms are complemented by an in-kind reservation of sovereign compute capacity and a separate, appropriately scaled mechanism addressing structural displacement within the creative sector beyond the function already performed by copyright licensing.

The case for a layered model follows from the unusually high degree of uncertainty surrounding the development of AI. There is presently no reliable basis for determining which technological and economic trajectory will prevail. The UK Government Office for Science (UK GO-Science) provides a useful framework through which to consider this uncertainty. Its AI 2030 scenarios encompass markedly different futures, ranging from a 'Slow Burn', in which technological progress slows and economic disruption remains relatively contained, to a 'Take-Off', in which AI systems outperform expert humans across almost all cognitive tasks within a comparatively short period. Australia, like other jurisdictions, cannot know with confidence which of these trajectories it is designing policy for. A policy architecture calibrated closely to one anticipated trajectory may consequently perform poorly under another.

AI policy therefore exhibits the characteristics of a VUCA environment, defined by volatility, uncertainty, complexity and ambiguity. This distinction has important implications for policy design. Proposals centred on a single mechanism necessarily embed assumptions, whether

explicit or implicit, about the future scale, structure and distribution of AI-generated value. Under conditions of deep uncertainty, the objective should not be to identify a single expected future and optimise an instrument around it, but to construct a portfolio of complementary mechanisms capable of performing reasonably across a range of plausible futures. Such a multi-instrument framework must, in turn, be calibrated carefully so that public participation in the economic value generated does not render investment commercially unviable or allocate risks and returns inequitably between participants.

The layered model is recommended on this basis. Its principal advantage is not that it maximises public revenue under any particular forecast, but that it provides a comparatively robust means of securing public value across substantially different technological and economic outcomes. If AI development slows, the more modest and infrastructure-based elements of the framework retain value. If AI instead produces substantial productivity gains and economic transformation, ownership and revenue participation allow the public return to scale with that success. The relevant policy criterion is therefore robustness rather than optimisation against a single forecast. Given the degree of technological uncertainty, reliance on any one instrument creates an avoidable concentration of policy risk.

The timing of these choices is equally important, and the absence of decisiveness may result in a limited return for Australians. The Prime Minister has announced that mandatory standards for AI and data centres will be legislated by early 2027 through a National Cabinet process. The announced framework includes requirements for developers to underwrite additional power supply so that costs are not transferred to households and businesses, and to return at least as much energy to the grid as their facilities consume. A new Office of AI within the Department of the Prime Minister and Cabinet will also oversee policy concerning AI, data centres and intellectual property.

Several elements that this report initially sought to establish as a policy baseline are therefore moving into government policy, providing an important foundation on which a broader value-capture framework can be constructed. The emerging baseline, however, does not by itself provide the public with an ongoing economic interest in the value generated through one of the largest prospective infrastructure programs in Australian history. Conditions governing entry and operation can determine how infrastructure is developed and the costs operators are required to internalise, but they do not necessarily provide the public with a continuing claim on the economic value generated by assets supported by Australian land, energy, infrastructure and institutional capacity.

This is the principal gap addressed by the recommendations in this report. Australia's bargaining position is arguably strongest before major capital commitments become irreversible and infrastructure, contractual arrangements and operating models become embedded. The experience of Australia's liquefied natural gas (LNG) sector provides a relevant domestic

precedent. Fiscal arrangements that proved comparatively weak in capturing subsequent economic rents became difficult to reform once investment had occurred and the associated infrastructure and commercial arrangements were established. The relevant lesson is not that the AI and LNG sectors are economically equivalent, but that value-capture arrangements are considerably more difficult to retrofit after large, geographically fixed investments have been committed.

3 Designing for an uncertain future

Any policy framework for AI training infrastructure must be designed to operate under conditions of substantial uncertainty. The defining feature of AI in 2026 is not any single technological capability, but the breadth of disagreement over the direction, pace and economic consequences of its development. This uncertainty is compounded by divergent assessments of the role and public value of open-weight models relative to proprietary systems, and by continuing uncertainty over where model development, training and deployment will ultimately concentrate. The UK Government Office for Science's *AI 2030 Scenarios*, developed with the AI Security Institute and updated in 2026, provides a particularly useful framework for considering these questions. Rather than advancing a single forecast, the framework presents five plausible scenarios structured around six critical uncertainties: the pace of capability development, the concentration of model development, the controllability of advanced systems, the extent of adoption, the degree of labour displacement, and the prospects for international cooperation or fragmentation. Their value lies precisely in their status as scenarios rather than predictions. They provide a basis for testing whether policy settings remain effective under materially different technological and economic conditions.

The five scenarios encompass a wide range of possible outcomes, with correspondingly different implications for policy design as shown in Table 1. These scenarios are not simply variations in the rate of technological progress. They describe materially different economic structures, labour-market outcomes and distributions of income, capital and economic power. They also alter the relative effectiveness of different policy instruments. A mechanism calibrated to capture a modest and relatively predictable stream of value may be adequate under 'Slow Burn' but comparatively insignificant under 'Take-Off'. Conversely, an intervention predicated on very large future rents may impose unnecessary costs if technological progress or commercial adoption proves more limited. The policy problem is therefore not simply to identify which scenario appears most probable, but to establish an institutional framework capable of producing acceptable outcomes across the scenario set. This is the basis on which the report assesses the case for a layered rather than singular approach to public participation and value capture.

Table 1: The five GO-Science AI 2030 scenarios, with two additional deployment-architecture scenarios

Scenario	Trajectory	What 2030 looks like	Implication for Australia
<i>Slow Burn</i>	Progress slows	AI outperforms humans at a minority of cognitive tasks; constrained adoption; modest uplift; displacement contained to a few roles; a market correction drags on growth	Even an underwhelming AI makes control of data and enabling infrastructure a durable advantage
<i>Open Frontier</i>	Progress slows	Open-weight systems narrow the gap to the frontier; many firms and nations compete; broad opportunity but widespread malicious-use capacity; China a leading supplier	Value and risk both diffuse; conditions and reserved capacity matter as activity spreads
<i>Augmented Growth</i>	Progress continues	AI matches most humans at most cognitive tasks; humans stay 'in the loop'; new jobs offset displacement; allies align on standards; a broadly shared boom	The benign case; all capture layers contribute and a growing sector can be built on
<i>Transformation Economy</i>	Progress continues	Similar capability, but humans pushed 'out of the loop'; severe displacement; rapid but uneven adoption; profits accrue overseas to a few foreign frontier firms	The future the model most guards against; value leaks offshore as labour income falls
<i>Take-Off</i>	Progress takes off	Systems surpass experts at almost all cognitive tasks within a few years; extraordinary growth; safety deprioritised amid race dynamics; a 'K-shaped' ⁴ economy	Value pools with owners of compute and capital; only ownership puts the public on that side
<i>Edge Computing</i>	Deployment shifts	More AI runs locally on devices rather than in central facilities; augmentation moves 'out of the loop' as on-device agents act with less oversight; inference disperses even as training stays concentrated	Capture anchored to deployment or output is harder to track as inference disperses; the case for anchoring to fixed training infrastructure strengthens
<i>Distributed Training</i>	Architecture shifts	Models become trainable across distributed, geographically dispersed compute rather than single large campuses; training	The fixed-asset leverage this report relies on weakens if training fragments; reserved compute, conditions and

⁴ A situation where different parts of the economy move in a different direction at the same time. Some sectors increasing in wealth, others stagnate or decline. This resembles the letter 'K'.

		itself becomes less reliant on any one site or jurisdiction	ownership of shared inputs (energy, grid) become the more durable anchors
--	--	---	---

Source: The first five scenarios are adapted from the UK Government Office for Science, AI Scenarios 2030 (June 2026); their names and characteristics derive from the GO-Science framework'. The Australian implications are based on Per Capita analyses. The final two scenarios, Edge Computing and Distributed Training, are deployment-architecture scenarios added by Per Capita to test the model against shifts in where and how AI runs rather than in capability alone.

3.1 Examining the common factors

While there are clear and unequivocal divergences, the GO-Science findings identify patterns that hold across most, if not all plausible futures, and these are the patterns that should serve as the basis of Australian policy. Three are central to the presented layered response framework. First, the frontier AI market is expected to remain highly concentrated, so that a large share of the gains accrues to a few frontier firms, to the owners of capital invested in them, and to those controlling key inputs; with rising inequality a recurring risk. Second, adoption spreads unevenly, so that the organisations, sectors and nations that capture disproportionate productivity gains pull away from those that do not. Third, global competition persists, led by the United States and China, notwithstanding the critical Mistral and Microsoft partnership in Europe, which will provide EU players with local infrastructure and access to Frontier grade models that are 'EU Made', so that outcomes for a country like Australia depend heavily on its access to technology, its partnerships, its position in the value chain and geopolitical positioning. In several scenarios the report is explicit that countries outside the frontier, risk seeing the value created on their soil flow offshore. Since the report, the EU has secured favourable outcomes through Microsoft/Mistral, leaving a number of middle powers such as Australia, currently without compute, frontier grade models, and less evidence of innovation and talent agglomerations.

These shared findings point to a single strategic problem for Australia. In most plausible futures, *value concentrates with those who own the frontier firms, the infrastructure and the scarce inputs, and a host country that merely facilitates activity captures little real value*. That is the problem that the layered approach directly addresses, the recommendations are designed to solve, and it is a problem that exists across the scenario range rather than only in the dramatic futures. Even in 'Slow Burn', where AI underperforms, control of high-quality data assets and enabling infrastructure becomes a source of durable economic and geopolitical advantage (particularly when it is a balanced pool of compute and inference assets). The case for capturing a public stake in AI futures does not depend on betting that AI will transform the world. With a balanced investment approach and layered intervention approach, it holds even under the 'Slow Burn' scenario.

3.2 A VUCA challenge calls for a balanced response

Forecasting under uncertain conditions is a challenge at the best of times. VUCA (volatility, uncertainty, complexity, ambiguity) is the standard description of decision environments where the range of outcomes is wide, the interactions between variables are hard to trace, and even the questions are contested. AI in 2026 fits the description, and may continue to be defined by VUCA conditions for some time. The GO-Science work makes the point in its own terms; it distinguishes uncertainty, which cannot be resolved by collecting more data, from risk, which can, and recommends scenarios precisely because the future of AI sits in the former category. Under genuine uncertainty, the instinct to identify the single most likely future and optimise for it, this is a flawed approach, some might argue a trap, because the cost of being wrong is borne in full when a different future arrives. Many proposals are being prepared with this limited mindset, that a single transaction tax, or sovereign royalty will resolve all challenges therein presenting a singular solution, but the mindset is unhelpful.

The disciplined and necessary alternative is to seek robustness rather than optimality and develop a policy that performs reasonably well across the whole range of plausible futures over one that performs best in the expected case and fails badly outside it. This is the principle that organises the recommendations that follow. For example, an access fee alone is well suited to a 'Slow Burn' world and badly suited to a 'Take-Off' one, where it would prove a trivial price for an enormous prize. A further example, an equity stake alone is powerful in the high-growth futures but ties up capital that looks expensive if AI disappoints. Each single instrument model proposed in public discussions of these issues, is a bet on a particular scenario. A layered model, by contrast, holds a position in each future at once, which is what robustness under uncertainty requires, a balanced approach that ensures a good outcome regardless of the future that eventuates. That robustness extends beyond how capable AI becomes, to how it is built and run. If inference disperses to local devices, or training itself fragments across distributed compute, a model that owns and conditions the shared inputs beneath the activity, still captures value, where one tied to a single training campus would not.

4 The opportunity and the window

Investment in AI infrastructure is accelerating faster than any comparable capital program in modern history.⁵ The decisions Australia makes in the next year will determine whether the public shares in the value created here, or simply hosts the activity and absorbs its environmental and social costs, as little more than a bystander. The government has signalled that it intends to negotiate the terms of AI training investment, rather than merely

⁵ Mohamed Hussein, 'Visualising AI Spending: How Does It Compare with History's Mega Projects?', *Al Jazeera*, 19 February 2026.

facilitate it. The Assistant Minister for Science, Technology and the Digital Economy has framed the choice as whether Australia shapes the boom before it arrives or reacts after it is here. This report takes that framing seriously and sets out what shaping the boom should mean in practice.

Training compute differs from inference in a way that is central to the policy design. Inference is distributed and mobile, it can run almost anywhere. Training is concentrated, capital-intensive and tied to specific corridors with the power, land and connectivity to support it. Hosting training therefore offers something hosting inference does not, that is economic value that is anchored to fixed assets within Australian jurisdiction. Assets that cannot be moved can be owned, conditioned and taxed in ways that mobile activity cannot. But this should not lead to the conclusion that co-investment and considered intervention are necessary in the context of compute assets alone, both must be subject to the same balanced and layered approach.

4.1 The policy window remains open, but leverage will shift

It is self-evident that the government has done some work in recent times to bring forward both the debate and needed policy and regulatory oversight, and over 2025 and 2026 the ground has moved quickly. The *Commonwealth's National AI Plan*, released in December 2025, set the strategic direction and flagged an intention to work with the States and Territories on national data centre principles⁶. In March 2026 the Australian Government followed it with its *Expectations of data centres and AI infrastructure developers*, a policy jointly announced by the Minister for Industry and Innovation, the Minister for Climate Change and Energy, and the Assistant Minister for Science, Technology and the Digital Economy⁷. Then, in July 2026, the Prime Minister announced that these Expectations would be legislated with mandatory standards for AI and data centres to be legislated by early 2027 through a national cabinet process, a new Office of AI will sit within the Department of the Prime Minister and Cabinet to oversee rules on AI, data centres and intellectual property, and data centre developers will be required to underwrite new energy generation. What began as a set of expectations framed around social licence is becoming an enforceable national framework, and this must be acknowledged and commended.

Several of what would, a year ago, have been the report's foundational or associated asks are now becoming the baseline. The new baseline established by government already asks operators to prioritise the national interest, to support the energy transition and cover a share of transmission and infrastructure costs, to use water responsibly, to invest in Australian skills and jobs through apprenticeships and structured training, and to strengthen local

⁶ Section 6 returns to this directly and stress-test the recommended model against all seven scenarios

⁷ Non frontier

capability by providing favourable compute access to start-ups, researchers and not-for-profits, and the legislation now foreshadowed will give several of these force.

The workforce and reserved-compute layers of this model therefore no longer need to be argued for as novel; the task instead is to specify how they should work and to ensure they are set at a level that genuinely serves the national interest. Our goal here is to provide a distinctive layered model that lies above this emerging floor. Conditions of entry, however strong, regulate how a facility behaves; they do not give the public an ongoing share in the value the facility helps create. Underwriting new power, returning energy to the grid and training apprentices are the price of a social licence to operate; ownership, revenue apportionment and a dedicated national fund are how the public comes to hold a durable stake.

The State processes point the same way. The New South Wales Legislative Council's Public Accountability and Works Committee is conducting a wide-ranging inquiry into the development and regulation of data centres, with a final report due by 30 September 2026⁸. Its terms of reference reach well beyond energy and water to the economic and distributional outcomes of the buildout: the distribution of costs and benefits, the extent of public subsidies and State facilitation, cost recovery for network and water infrastructure, and workforce development and training. The committee has signalled that it will consider recommending specific financial contributions from data centre proponents toward infrastructure costs. In parallel, the New South Wales Government has issued a consultation paper canvassing whole-of-government approaches to regulating the sector, and the Australian Energy Market Commission has released draft connection reforms that would require large electricity users, including data centres, to remain connected and help stabilise the grid during disturbances. The distributional question this report addresses, who captures the value created when AI is trained on Australian soil, is therefore already live in at least one State parliament and across several Commonwealth processes at once.

4.2 Why the window is closing

Australia's leverage rests on advantages that are most valuable before infrastructure is committed. Renewable energy access, rule-of-law stability, accelerated approvals and a robust copyright framework are bargaining chips now, but these advantages will be carried only for a short term. Once campuses are built, grid connections are made and supply chains are established, the leverage passes to the firms that operate them. This is not speculation, this has been observed in other jurisdictions and it is the documented pattern of Australia's offshore gas development.

⁸ New South Wales Legislative Council, *Portfolio Committee No 7 – Planning and Environment, Inquiry into Data Centres in New South Wales*, 2026.

Despite being among the largest LNG exporters in the world, Australia has in some years collected close to zero in Petroleum Resource Rent Tax (PRRT) from the sector, because deductible capital expenditure, uplift provisions and project-level accounting interacted to defer or erase the rent the tax was meant to capture. The Henry Tax Review and the Parliamentary Budget Office have both documented the structural gap between the country's resource wealth and its fiscal return. Treasury has acknowledged that the PRRT was not designed for the capital intensity of LNG. The lesson is not that resource taxation is impossible. It is that an instrument must be designed for the specific economics of the industry in question, reflecting different trajectories rather than being based in a single assumption; and designed before the leverage is lost. A framework built for data centre construction, will not capture the national interest in AI training if it ignores how AI models are actually developed, deployed and monetised. This is why the layered framework is essential.

5 What other jurisdictions are doing

At the time of writing, no jurisdiction has constructed a comprehensive value-capture regime for AI training infrastructure. Activity accelerated markedly through the first half of 2026, but the international position remains piecemeal, and each initiative that has emerged rests upon a single instrument and, by implication, upon a single view of how the technology will develop. That fragmentation is the reason the design window remains open.

Table 2: Emerging approaches in other jurisdictions

Jurisdiction	Proposal or approach	Instrument base	Status and issues
United States, federal	A series of federal proposals has considered alternative mechanisms for capturing value associated with AI and data centre development; the <i>American A.I. Sovereign Wealth Fund Act</i> (Sanders, 18 June 2026); <i>AI Tax and Work Protection Act</i> , H.R. 10044 (Casar, 6 August 2026); Wyden's <i>Data Center Public Investment</i> excise-tax white paper (6 August 2026); and Warren's proposed data centre energy excise (advocated May 2026, with no bill introduced).	Company equity; token use or AI-derived revenue; data centre gross receipts; energy consumption	The four proposals advanced within an eleven-week period illustrate both growing congressional interest in AI-related value capture and the absence of consensus regarding the appropriate tax or revenue base. None appears to have a clear legislative pathway in the current Congress. H.R. 10044 was referred to the Committees on Education and Workforce and Ways and Means, while consultation on the Wyden proposal remained open until 31 August 2026.
United States, states	State-level responses have principally involved the repeal, suspension or conditioning of existing sales and use tax exemptions for data centres, rather than the creation of new AI-specific revenue instruments.	Withdrawal or conditioning of existing fiscal incentives	This represents the most developed area of implemented fiscal intervention identified across the jurisdictions considered. At least twenty-five states introduced legislation during the 2026 sessions to repeal or restrict data centre tax concessions, with eight measures enacted. Maine repealed relevant subsidies, Arizona introduced a three-year moratorium extending to June 2029, and North Carolina repealed its electricity exemption. Illinois, Massachusetts and Ohio also

PER CAPITA – AUSTRALIA'S AI DIVIDEND

			paused incentive programs through executive or administrative action.
South Korea	A “national dividend” was proposed by the presidential chief of staff for policy on 12 May 2026, drawing on additional tax revenue associated with growth in the AI and semiconductor sectors. Proposed applications included youth start-up capital, rural basic income, support for artists and pension provision.	Incremental or excess tax revenue rather than a dedicated corporate or AI levy	The proposal should be distinguished from formal government policy. The presidential office subsequently stated that it had not been subject to formal policy discussions, while the President clarified that the mechanism would not be financed directly from corporate profits. The initial announcement was associated with a significant adverse market response. Its immediate incidence was principally relevant to semiconductor and memory-chip producers rather than AI developers. Separately, legislation remains pending for a sovereign wealth fund exceeding 20 trillion won for strategic industries, financed through budgetary allocations rather than AI-specific levies.
European Union	The AI Act establishes an extensive compliance framework but does not provide a mechanism for capturing AI-generated revenues. Separately, digital services, gambling and crypto-related levies are being considered as potential own resources within the 2028 to 2034 EU budget framework.	Regulatory compliance obligations; digital-services and related revenues	The European approach remains primarily oriented towards risk regulation rather than the fiscal capture of AI-generated value. The European Commission estimates that the proposed own-resources package could generate up to €11 billion annually, although adoption requires unanimity among Member States. The AI Liability Directive was formally withdrawn in October 2025, removing the previously proposed liability pathway. The development of digital levies also occurs within the context of potential United States trade responses.
United Kingdom	The <i>AI Opportunities Action Plan</i> (January 2025) emphasises infrastructure investment, regulatory reform and the development of domestic AI capacity. The Government's copyright and AI consultation subsequently concluded with a report published on 18 March 2026.	Investment and regulatory facilitation; no dedicated revenue-capture instrument	The UK approach remains predominantly investment-oriented and does not propose an AI-specific sovereign fund or comparable revenue-capture mechanism. The Government did not proceed with its previously preferred model of a text and data mining exception accompanied by rights-holder opt-out, instead indicating that further evidence would be gathered. The UK experience is relevant to Australia because it illustrates the policy tensions between attracting AI investment and maintaining established copyright protections in a comparable legal and institutional setting.
India	India has adopted a tax holiday extending to 2047 for revenue from cloud services supplied outside India where the underlying workloads are processed in Indian data centres, alongside a 15 per cent cost-plus safe harbour for domestic operators providing services to related foreign entities.	Fiscal incentive rather than revenue capture	India represents a contrasting policy strategy in which fiscal settings are being used to attract internationally mobile data centre investment rather than to capture a greater share of associated rents. Announced through the annual budget, the approach is relevant to Australia because it demonstrates the international competition for AI and cloud infrastructure and, consequently, the external constraints that may affect the terms host jurisdictions can negotiate.

The instructive precedents lie outside AI, in the fiscal architecture built around concentrated, externally-driven revenue streams (see Figure 4). Norway's Government Pension Fund Global, built from North Sea oil and gas and now the world's largest sovereign wealth fund at more than US\$2 trillion (around A\$3 trillion)⁹, embedded three features Australia lacked in its LNG arrangements; equity participation rather than taxation alone, a clear separation between

⁹ Ritsuko Shimizu, 'Norway Sovereign Wealth Fund', *CNBC*, 27 February 2026.

resource revenues and the operating budget, and an offshore investment mandate to avoid domestic inflationary effects. Chile's copper-funded stabilisation fund shows how to insulate a budget against a volatile, concentrated revenue stream. Ireland's Future Ireland Fund, capturing windfall corporate-tax receipts from a foreign-dominated sector, is the closest contemporary parallel to capturing hyper-scaler rents. The transferable lesson is the architecture, not the resource, and Australia need not build from scratch; the Future Fund already provides a credible vehicle.

Figure 4 – The design of international sovereign funds

Institutional Feature	 Norway Government Pension Fund Global	 Ireland Future Ireland Fund	 Chile Economic and Social Stabilisation Fund	 Australia Proposed Australian AI Dividend Fund
Revenue separated from operating budget	● ● ● Separate from central government budget	● ● ● Statutory fund outside annual budget	● ● ● Ring-fenced from fiscal spending	● ● ● Quarantined from operating budget
Long-term investment mandate	● ● ● Very long-term horizon, diversified globally	● ● ● Intergenerational investment horizon	● ● ○ Medium to long-term mandate	● ● ● Long-term horizon to build national wealth
Manages volatile or external revenue	● ● ○ Oil revenues are cyclical; buffered over time	● ● ○ Designed for windfall and volatile receipts	● ● ● Explicit stabilisation mechanism for volatility	● ● ○ Mandate to manage externally-driven volatility
Intergenerational orientation	● ● ● Strong intergenerational principle	● ● ● Built explicitly for future generations	● ● ○ Supports long-term fiscal sustainability	● ● ● Benefits current and future Australians
Independent governance	● ● ● Strongly independent from political influence	● ● ○ Governed independently with clear mandate	● ● ○ Independent committee oversight	● ● ● Independent governance for durability
Direct link to concentrated revenue source	● ● ● Directly linked to petroleum reserves	● ● ● Linked to corporate tax and resource receipts	● ● ● Linked to copper revenues and fiscal rules	● ● ● Directly linked to AI infrastructure revenues

● ● ● Strong
● ● Moderate
● Limited

Note: This is an overview of different sovereign fund models, based on Per Capita analysis.

6 Why a single instrument fails

Australia's existing corporate tax framework is not an adequate mechanism for capturing value from AI training. Major AI companies are structured to minimise taxable profit in higher-taxing jurisdictions through intellectual property held in low-tax entities, intercompany loans and

cost-sharing arrangements that shift value offshore. Therefore, any design that relies on taxing accounting profit will repeat the PRRT outcome.

That rules out the easy option and forces a more deliberate choice among instruments that each capture value in a different way and fail in a different way.

Table 3: Capturing value from AI infrastructure deployment and use

Instrument	What it captures	Characteristic weakness
Access fee	Recognises the public endowment a facility consumes. Land, transmission, clean water, greenfield sites and a stable grid; simple to administer; not a copyright charge. This fee is over and above transmission and grid contribution expectations outlined by government.	Essentially a one-time charge on entry; captures the access but not the value subsequently created; hard to renegotiate later.
Token tax	Ongoing value from AI usage by volume; cannot be minimised through transfer pricing	Yield depends on how broadly the taxable activity can be defined and enforced
Revenue participation	An ongoing share scaling with the commercial value of models that benefited from Australian inputs	Defining 'materially benefited' invites the same deduction and transfer-pricing pressures as the PRRT unless anchored to countable units
Co-investment	A direct equity share of the training asset; most resistant to profit-shifting; closest to the Norwegian model	Significant negotiating and diplomatic resistance; capital commitment
Infrastructure ownership	A public stake in the fixed enabling assets, which are geographically anchored and long-lived	Capital intensity and the timing of entry; far cheaper before commitment than after
Reserved sovereign compute	An in-kind share of capacity for public-interest use; attaches to a physical asset in jurisdiction	Secures capability rather than fiscal return; complements rather than substitutes for revenue mechanisms
Training and apprenticeship contribution	Captures value as human capital. Trained apprentices and a deeper skilled-trades workforce; attaches to the worksite, which cannot be relocated	Builds workforce capacity rather than fiscal return; effective only if matched by public investment in training places, and best set in dialogue with the training sector

The pattern emerges from the table, the instruments that are easiest to administer capture the least lasting value, the smallest dividend for the community. The instruments that capture the most durable value are the hardest to negotiate.

For example, an access fee is clean but shallow, in contrast co-investment and infrastructure ownership are deep but capital-intensive and diplomatically contested. Revenue participation scales with success but is only as good as its anti-avoidance design. However, that is not the reason to choose the weakest option. It is the reason to combine them.

7 The recommended model: layered capture through partnering and shared ownership

The recommendation is a layered model in which public-private partnering and shared ownership of the fixed assets form the structural core, an ongoing revenue apportionment captures value as it is realised, and in-kind reservations secure sovereign capability and a skilled workforce. Each layer is embedded so that its strength corresponds to the other layers' weakness. The access fee provides an immediate floor; ownership provides durability; apportionment provides scale; reserved compute provides strategic capability; and the training contribution builds the workforce the buildout depends on to be completed. The layers are mutually reinforcing (rather than alternatives), and the model degrades gracefully by design, if one layer is negotiated down, the others still hold.

7.1 Layer one: shared ownership within a partnership framework

The first layer is the structural core of the model, a durable public stake in the fixed enabling assets, held not as a passive holding but inside a genuine partnership with governance rights. It has two components, set out below, that work together. The equity stake gives the public a claim that mobile activity and accounting profit cannot erode. The partnership framework turns that claim into influence over how the asset is run. For emphasis this captures both compute and inference assets, and this is a source of balance and resilience.

Shared ownership of the fixed assets

The foundation is a public equity stake in the enabling assets themselves, the data centre campuses, compute facilities, transmission infrastructure and the renewable generation behind them. These assets are geographically fixed, long-lived and concentrated in identifiable corridors. That is exactly what makes ownership durable in a way that taxing mobile activity is not. In many industries, owning the infrastructure has proven a more stable source of public benefit than taxing what is done with it, because the asset cannot be relocated to a more favourable jurisdiction when the arrangement becomes inconvenient.

Australia already holds such stakes in a number of assets through established vehicles. For example, the Clean Energy Finance Corporation takes equity in energy and enabling infrastructure as a matter of ordinary practice. Snowy Hydro and NBN Co are precedents for

public ownership of nationally significant infrastructure. The question is therefore not whether the state can hold these stakes but whether it extends an existing instrument/vehicle to a new asset class. The binding constraint is timing, it is critical that decisions are made with some expediency. Stakes are far cheaper to negotiate before infrastructure is committed than after. This may feel like a common refrain to this point but emphasis is necessary.

A genuine partnership, not passive subsidy

Shared ownership must be structured as a genuine partnership rather than a grant presented as partnership. The distinction is governance, and maintaining some degree of ethical influence over asset use. A public stake should carry board representation, information rights, transparency over training runs and model deployment, and step-in rights in defined circumstances. Public co-investment should be conditioned on the partner accepting these terms, and on compliance with the Australian Government's Expectations of data centres and AI infrastructure developers, which already set out principles covering national interest, energy, water, skills and research access.

The reason to insist on partnership terms is that passive concessions reproduce the LNG failure in a new form. This is simply not an option given the deleterious impacts of this faulty earlier model. A tax concession or an unconditioned grant transfers public value to a private operator and retains no claim on the upside and no seat at the table when decisions are made. A partnership retains both, affording the community an enduring benefit and enduring influence. It also gives Australia visibility into activity that is otherwise opaque, which is itself a national-interest benefit, given the transparency and data-sovereignty questions the training of frontier models raises. The impact on the technologists is modest, moderately lower revenue and ethical oversight are a foundational baseline.

7.2 Layer two: revenue apportionment anchored to countable units

Ownership captures a share of the asset, but it does not by itself track directly the commercial value the asset helps create downstream. That is the role of revenue apportionment, an ongoing ad valorem share tied to the commercial utilisation of models that have materially benefited from Australian infrastructure, energy, public investment or copyright protections. Where a token tax charges by volume, apportionment charges by value, so it scales with commercial success rather than raw throughput, and it is harder to escape by relocating training, because the trigger is deployment revenue rather than the training event. This layer does not sit onerously with the former as both may be calibrated to ensure fair and commercial viable rewards for all parties.

The risk is largely administrative, and it is the same risk that eroded the PRRT. Defining 'materially benefited' and the relevant revenue base invites transfer-pricing and deduction pressure. The design response is to avoid taxing accounting revenue at all. Anchoring the

charge to discrete, countable licence units, such as registered models or deployment instances, gives the anti-avoidance architecture something solid to attach to, because a per-unit charge cannot be minimised through transfer pricing in the way a profit-based charge can. So framed, revenue participation borrows the enforceability of a token tax while still scaling with commercial deployment. It remains only as good as its anti-avoidance design, but anchoring it to units is what makes that design tractable.

Flatter mechanisms are plausible, but they increase the risk that the public share of the value created will be comparatively modest, particularly where the commercial value generated by AI grows substantially over time. A token tax provides one alternative, imposing a charge on the volume of tokens processed or generated by a model. Such an approach has considerable administrative appeal because tokens provide a relatively observable and countable unit of AI activity, and a volume-based charge would be less susceptible to conventional profit shifting and transfer-pricing strategies. Its weakness, however, is that the tax base is closely connected to the marginal economics of AI use itself. If set at a material rate, a token tax could increase the cost of inference and potentially alter decisions about model use, efficiency, architecture and where workloads are processed. If set sufficiently low to avoid materially affecting those decisions, its capacity to capture a meaningful share of the potentially very large economic rents generated by successful frontier models may be correspondingly limited. The policy challenge is therefore not simply whether tokens can be taxed, but whether they provide an appropriate basis for capturing economic value without materially distorting the underlying economics of AI production and deployment. Where such balance is achievable a token tax is an equally plausible alternative.

Nonetheless, the choice of countable (license) units is a pragmatic proxy, adopted because the thing the apportionment is really trying to track, the contribution that Australian inputs make to a model's commercial value, cannot yet be measured directly. That underlying idea (not the basis of measurement at this point) is worth naming, because it is what gives the apportionment its principled justification rather than leaving it as an arbitrary charge. It is based on what this report calls the Attributable Influence Principle, that where a work or an input makes a measurable causal contribution to the productive capability of an AI system, that contribution is a legitimate basis for a share in the value the system generates, even when nothing is reproduced verbatim. Reproduction is the wrong test for a technology that transforms its inputs into distributed statistical representations; contribution is the right approach, albeit one that is offered as foundational conjecture at this point. The presented option, aligned to apportionment (rather than attribution), is countable-unit-based today as a matter of enforceability, but the principle it serves is contribution, and as the attribution science matures the same charge could in time be calibrated to measured influence rather than to a proxy. Appendix A develops this principle in full and sets out the emerging techniques that make it increasingly tractable. Notably the apportionment method could be

simplified to account for a mutually agreed revenue stream apportionment, where accounting sophistry is not allowed to offset its determination, as was the case with LNG.

7.3 Layer three: an access fee recognising the public endowment

Beneath the noted layers in the model sits an access fee that provides an immediate revenue floor. It is worth being precise about what this fee recognises, because it is not a charge on training as an activity and still less a payment for the works a model learns from. It is recognition of the public endowment that makes hosting training here possible at all, the land and greenfield sites that large facilities consume, the transmission capacity and the firmed, increasingly clean electricity drawn from a shared grid, the clean water used for cooling, the accelerated planning and approvals, and the stability of a well-governed jurisdiction. These are finite, publicly stewarded resources. When a hyperscale facility takes a grid connection, a parcel of industrial land and a water allocation, it is drawing on an endowment that belongs in part to the public and to future users who will compete for the same scarce inputs. The access fee is the price of that draw, set in recognition of the opportunity Australia extends rather than as a tax on any particular output.

In this way, the fee is clean to administer and conceptually distinct from copyright. It attaches to physical access, which is observable, sited and hard to relocate, so it avoids the unworkable idea of policing or pricing the act of training itself, an idea that would be both implausible to administer and an intrusion into territory that Australia's existing copyright provisions already govern. The fee carries a second rationale beyond resource access. As Appendix B sets out, the deployment of models trained here draws on a finite pool of human attention and can erode the economic base that sustains the journalism, research, education and creative work on which a healthy knowledge economy, and indeed future AI capability, depends. A modest, broad-based access fee is a reasonable way to recognise that wider public interest in sustaining human knowledge production, alongside the physical endowment the facility consumes. The fee should not be set without reference to long-term upside, because a fee fixed in isolation is a one-time charge whose inadequacy only becomes visible in retrospect. In the layered model it does not have to carry that weight, because the ownership and apportionment layers capture the upside it cannot. This fee could also be offset against future revenue capture should future revenues achieve a level of sufficiency. It is also plausible that the fee could serve as an annual minimum charge that is offset against throughput revenue and therein serving as a revenue floor, ensure no repeat of the LNG scenario, though this approach may be unnecessary with a well designed second layer.

7.4 Layer four: reserved sovereign compute

The fourth layer captures value in kind rather than in monetary terms. Access to Australian infrastructure should be conditioned on a reserved share of compute capacity for universities,

public-interest research, start-ups, the national science agency and strategic national needs. The model is Western Australia's domestic gas reservation policy, which has proven durable precisely because the reserved share attaches to an asset that cannot be moved. The benefit here is strategic capability rather than revenue, which is why reservation complements the cash mechanisms rather than replacing them. It is worth noting that Australia's principal national science agency has been significantly depleted, so a reservation that includes uplift of that agency (through monetary or compute allocations) would serve a national interest beyond the AI sector itself.

7.5 Layer five: a training and apprenticeship contribution

The final layer captures value as human capital, that must be seen as central to the success of this rollout. AI training infrastructure is, in physical terms, a very large electrical project drawing on a skilled-trades workforce that is already scarce, so the buildout creates an acute labour pressure at the same time as it creates a rare opportunity to train new entrants. This layer turns that pressure into a lasting public benefit, creating a once in a generation opportunity for sectoral uplift, and serving as an uplift rather than a hinderance to the energy transition. This is achieved by attaching a training and apprenticeship commitment to the same access that the other layers price, operators benefiting from Australian infrastructure contribute to developing the local skilled workforce, whether through a set apprentice-to-tradesperson ratio, or a contribution to a training fund that supports site-based apprenticeships and the surrounding training ecosystem. Like reserved compute, the value here is captured in kind rather than in monetary terms, and like the access fee it attaches to a fixed, sited asset that cannot be relocated, which is what makes the commitment enforceable. Because data centre sites are among the largest, cleanest and safest electrical worksites in the country, and most sit in metropolitan areas close to existing training infrastructure, they are unusually well suited to carrying it. The mechanism, and the verified workforce evidence behind it, is set out in full in Section 9, where this layer also operates as a national-interest condition. The settings should be developed in dialogue with training entities, the TAFE sector, group training organisations and unions, and matched by public investment in training capacity. The opportunity to uplift the TAFE sector and in turn see cascading benefits to the Housing sector and the Energy sector are self evident but warrant being noted. The scope for significant spillover effects to other sectors is more probable than merely plausible.

8 Stress-testing the model against the 2030 scenarios

The merit of a layered model is that it does not depend on guessing which future arrives, and basing the response around a singular assumption. The recommendations are tested directly against all seven scenarios from Table 1, asking of each, in this world, does the model still

capture value for the public, and which layers do the work? The five capability scenarios test how much value there is to capture and how it is distributed. Then the two deployment-architecture scenarios test something different and arguably much harder, whether the value stays capturable when the activity the model is anchored to begins to disperse.

The point of the exercise is not to show that every layer fires in every future. It is to show that in each future at least one layer captures meaningful (public) value, so that the Australians are not left empty-handed whichever way AI develops. That is what robustness under uncertainty looks like in practice, and what makes a layered approach the most viable. For emphasis it is pertinent to note that the private data centre provider is treated fairly under this model as well as there is no single value capture mechanism that is so onerous that it makes the underlying compute and inference economics unviable. It is a mix of well calibrated mechanisms rather than a single blunt instrument.

Table 4: How the layered model performs across the seven scenarios

Scenario	What this world looks like	How the layered model performs
<i>Slow Burn</i>	Progress stalls; modest uplift; displacement contained; control of data and infrastructure still confers advantage	The access fee and reserved compute carry the load and look well-judged; equity stakes are inexpensive to hold and the apportionment yields very little but costs are commensurate. Australia captures a fair price for access without having over-committed capital
<i>Open Frontier</i>	Open-weight systems narrow the gap; many firms and nations compete; opportunity and malicious-use risk both spread	Reserved compute and national-interest conditions become especially valuable as activity diffuses; the access fee still prices entry; ownership of fixed local assets anchors value even as software competition intensifies
<i>Augmented Growth</i>	Continued progress; humans 'in the loop'; broadly shared economic boom; allies align on standards	All layers contribute to the return; equity stakes appreciate, apportionment scales with strong commercial deployment, the fund compounds, and reserved compute builds sovereign research capacity into a growing sector
<i>Transformation Economy</i>	Severe displacement; humans 'out of the loop'; profits leak offshore to a few foreign frontier firms; fiscal strain	This is the future the model most guards against; ownership and apportionment recapture value that would otherwise leak overseas, and the fund plus the creative-sector mechanism provide a domestic claim on offshore-bound profit precisely when labour income is falling
<i>Take-Off</i>	Systems surpass experts at almost all cognitive tasks; extraordinary growth; value	Equity in the fixed assets and the apportionment become the decisive layers, placing the public on the capital-owning side of the divide; an access fee

	pools with owners of compute and capital in a 'K-shaped' economy	alone would have been a trivial price for an enormous prize, which is exactly why it is not relied on alone
Edge Computing	Inference disperses onto local devices and on-device agents act with less oversight, even as training stays concentrated in large facilities	Because training remains sited, the ownership, access fee and reserved-compute layers are unaffected; the apportionment is the exposed layer, since deployment becomes harder to count, which is an argument for anchoring it to registered models rather than to dispersed inference. The model holds, with the apportionment recalibrated
Distributed Training	Training itself fragments across geographically dispersed compute, so it is no longer tied to a single large campus or jurisdiction	This is the scenario that most tests the model. The layers tied to a single training site weaken, but the layers tied to shared, immovable inputs do not: the energy and grid the dispersed nodes still draw on, the reserved compute and conditions attached to any Australian capacity, and ownership of transmission and generation. Capture shifts from the campus to the inputs beneath it

Two key observations follow from Table 4, pertaining to the balance between conservatism and ambition. First, the instruments/mechanisms that look most ambitious and hardest to negotiate, ownership and apportionment, are the layers/mechanisms that matter most in the futures where the stakes are highest, namely 'Transformation Economy' and 'Take-Off'. A cautious regime that omits them performs adequately only if AI underwhelms, which is to say it is a concentrated bet on the 'Slow Burn' corner of the scenario space. Second, the simpler instruments are not redundant. The access fee and reserved compute are the layers that justify themselves even if AI disappoints, and they provide the immediate, legally clean floor on which the more ambitious layers are built. The model is deliberately over-specified relative to any single future, because that redundancy is the source of its robustness.

A third observation follows from the two architecture scenarios, and it is worth stating plainly because it identifies the model's real point of exposure. The capability scenarios test how much value there is to capture; the architecture scenarios test whether the value stays capturable at all. 'Edge Computing' is the milder case; training stays sited (or largely sited given progress in distributed training), so only the apportionment is stressed, and anchoring it to registered models rather than to dispersed inference largely answers the problem. Notably balance with regard to relative ownership of inference and compute assets, rather than compute in isolation also supports more balanced and assured outcomes in each scenario notwithstanding other factors. 'Distributed Training' is the harder case, because it erodes the single premise the model leans on most, that training is geographically anchored. The honest conclusion is not that the model fails in that world but that its centre of gravity shifts, capture moves from the training campus to the shared inputs the dispersed nodes still cannot do without, namely the energy, the grid connection and any reserved Australian

compute. And again, for emphasis a further reason why there must be balance in investment in inference and compute infrastructure. This is a further reason the model is layered rather than singular. A regime built only on owning the (compute infrastructure) campus would be fragile to distributed training; a regime that also owns and conditions the inputs beneath the campus is not.

This is also why the sequencing in Section 2 matters. The access fee and reserved compute can be settled quickly and will earn their place in every scenario. The ownership stakes and the apportionment, which are decisive in the high-impact futures, are the ones that must be negotiated while leverage is at its peak, because they are the layers that become impossible to retrofit once the infrastructure is built. Designing for uncertainty and designing before the window closes are, in the end, the same instruction/intention, an acknowledgement of the VUCA conditions.

9 Governing the proceeds

Capturing value is only half the task, the other half is ensuring the value is treated as a long-term national asset rather than a short-term income stream to be spent in the political cycle in which it arrives. This is the structural lesson of Norway's fund and the structural failure that Australia has repeated more than once. The cash proceeds of the access fee and revenue apportionment, and any dividends from the equity stakes, should be quarantined in a dedicated fund on the fiscal architecture of the (current) Future Fund or, modelled on Ireland's well regarded Future Ireland Fund.

Three design features are important, each of which is highly consequential. First, a clear separation between the captured revenues and the operating budget, to prevent the long-term asset value being spent as current income. Second, an investment mandate that manages the externally-driven and potentially cyclical nature of these rents, which is the lesson of Chile's stabilisation fund if AI infrastructure rents prove volatile. Third, governance independent enough to survive changes of government and commodity-style price cycles, which is the feature that made the Norwegian arrangement durable across decades. Australia does not need to invent this architecture. It needs to point an existing vehicle at a new revenue source.

Land value capture is a useful front-end complement at the state level. Rezoning, grid connection and fast-tracked approval create a land windfall through the planning act itself, capturable at approval through betterment levies or the ACT's Lease Variation Charge. It is one-time and site-bound, so it pairs naturally with the ongoing revenue and ownership mechanisms, and because it is a state lever the value accrues outside any Commonwealth fund. It is noted here for state-level consideration rather than as part of the core

Commonwealth model. The work of Prosper Australia offers significant insight into plausible approaches to land value capture.

10 Protecting the creative sector

Australia's copyright framework is the source of much of the leverage described in this report, and it is important to be clear about what that framework does, including where it is now changing. Compensation for the use of creators' work is the province of copyright (and not the focus of this report), and this is an area the Government has signalled it intends to strengthen, which must be commended. In July 2026, the Prime Minister stated that no company should use Australian books, music, art or news to train AI without the artist's control, including control over the price and value of their work, and the Attorney-General is facilitating consultation on copyright and artist protections where AI training is involved. This report does not propose to replace, supplement or second-guess that system, whatever form the strengthened settings ultimately take. The question it raises here is a different one that copyright, current or reformed, was never built to answer, whether the creative economy can remain viable as a whole when AI compresses the market for creative labour. It is useful to separate three problems that current debate tends to conflate:

1. Securing consent, control and payment where creators' work is used to train models, which is the province of copyright and of the reforms now under consideration;
2. The ongoing displacement of creative income as AI-generated content substitutes for commissioned human work; and
3. Whether the current economic model can sustain professional creative practice at all over time.

Existing copyright licensing addresses the first problem, but copyright was not designed to deal with all problems, not should it be. The second and third are problems of structural economic displacement, and they fall outside what any licensing scheme was designed to do. The Public Lending Right scheme is the precedent most often cited for a displacement response, and it has administrative merit, but it was designed to compensate for reduced sales caused by library circulation, not to sustain an industry under structural threat. Its payments are modest and its eligibility criteria were not designed for the AI context. A response calibrated to the actual scale of disruption would need to be more ambitious, more independent and more adaptable, and it would need to reach beyond named creators to the broader cultural economy of journalists, designers, translators and others whose markets AI is restructuring. The recommendation is therefore to design a creator-support mechanism for ongoing structural displacement, distinct from and additional to copyright licensing, funded in part from the value-capture proceeds, so that the sector most exposed to AI shares directly in the dividend its work helps generate.

How such a mechanism directs support is a question that attribution science may, in time, help answer. As noted, Appendix A sets out the Attributable Influence Principle, a forward-looking account of how measurable contribution to a model's capabilities might eventually be used to calibrate both the revenue apportionment and this creative-sector mechanism. It is acknowledged that much of the literature has centred around other forms of Machine Learning and Deep Learning and the field is an emergent one, but the shared literature provides an evidentiary basis for further exploration, something that should be encouraged. It is included as a horizon for the model rather than a precondition for acting, what is proposed within the appendices does not preclude acting now with a simpler approach. Notably, it does not disturb the division of labour set out above, in which licensing remains the province of existing copyright provisions.

11 Conditions on the training privilege

Several of the conditions discussed here are moving from aspiration to law. The Australian Government's July 2026 announcement that data centre and AI standards will be legislated by early 2027, strengthening the March 2026 Expectations, means that requirements on energy, water and skills are becoming a mandatory baseline rather than a matter for negotiation. That is welcome, and it changes the entire nature of the report process from a baseline standpoint. The task is no longer to argue that conditions should exist, it is to ensure they reach beyond the operational floor to the national interest in value capture. Training in Australia should comply with Australian law, but compliance is the floor, not the ceiling. The privilege of training on Australia's advantages should carry conditions that reflect the national interest. The available conditions beyond baseline environmental and planning requirements include notification of training runs, transparency about model architecture and training-data composition, data-sovereignty requirements about where Australian data is stored and processed, prohibitions on uses contrary to Australian law or national security, a commitment to developing local talent and the skilled trades workforce, and requirements to contribute to domestic AI research capacity.


On consent, a more complex legal and moral question in many ways, the live question is opt-out versus opt-in. Rights-holder groups argue for opt-in or, at minimum, an opt-out that is technically simple, free and effective across all types of works. The layered ownership model strengthens Australia's position in this context, because conditions attached to an asset the public co-owns and to compute the public has reserved are more enforceable than conditions attached to mobile, distributed activity. The hardest challenge is enforceability against large global firms with limited Australian legal presence, including the risk of diplomatic pressure under Section 301 of the United States Trade Act. Conditions must be designed to withstand

that pressure, which is a further argument for anchoring the regime to fixed assets and countable units within jurisdiction rather than to activity that can be relocated or restructured.

11.1 A commitment to local talent and apprentices

One national-interest condition deserves particular emphasis, because it converts a constraint on the buildout into a lasting public benefit. AI training infrastructure is, in physical terms, a very large electrical project, the data centres, the compute halls, the substations and the grid connections that power them require a substantial skilled-trades workforce, electricians foremost among them. That workforce is already scarce. Jobs and Skills Australia's 2023 Clean Energy Generation study found that electricians are the single largest area of workforce shortage among the critical occupations for the energy transition, with Australia needing an estimated 32,000 additional electricians by 2030 and around 85,000 by 2050. That study was completed before the current surge in data centre proposals, so the electrical demand created by AI infrastructure sits on top of an already-stretched pipeline rather than within it. Left unmanaged, a wave of large data centre builds competing for the same trades, could slow not only its own delivery, but also the housing and renewable-energy programs drawing on the same workforce.

Figure 5 - Electrical systems within data centres and at points of connection



IEA. CC BY 4.0

Note: a foundational overview of core systems based on IEA 2026 diagrams and analysis.

This could also be seen as a unique opportunity to advance all three sectors, if the government embeds simple, and largely cost effective (for both firms and operators) training mandates. This is a constructive response to a unique challenge. Because a data centre cannot be built without a grid connection and planning approval, those approval points are a natural

place to attach a commitment to developing local talent. The condition should be a genuine training obligation, not a box-ticking exercise, operators benefiting from access to Australian infrastructure should commit to train and employ Australian apprentices and to invest in the local skills base, including on the renewable-energy projects connected to their facilities. Data centre sites are well suited to this as they are among the largest, cleanest and safest electrical worksites in the country, most are in metropolitan areas close to existing training infrastructure, and their scale offers apprentice opportunities that few other projects can match.

One workable mechanism is an apprentice-to-tradesperson ratio; for example, one apprentice for every five qualified trades workers, with particular attention to first and second-year apprentices so that the obligation creates genuinely new entrants rather than drawing experienced apprentices away from small businesses. The ratio option is worth serious consideration, given the logical utility to all parties. The right settings/calibration, including the precise ratio, how it is verified, and whether shortfalls are addressed through a training contribution or other means, should be developed in close dialogue with training entities, the TAFE sector, group training organisations and unions, because those bodies hold the practical knowledge of where training capacity exists and where it must be built. A commitment that outruns the system's capacity to deliver training places would not produce more electricians; it would simply lengthen the existing queue. Any obligation on operators should therefore be matched by public investment in training capacity, so that the demand the sector creates can actually be met. Designed this way, in true mutually beneficial partnership, a local-talent commitment turns the significant workforce pressure that data centres create into one of the strongest trade-training opportunities the economy has seen.

12 Trade-offs and the case for ambition

The recommendations in this report sit toward the more ambitious end of the available spectrum. That is a deliberate response to the timing argument, not an oversight of the risks, the risks are well acknowledge but so to is the opportunity, and the layered approach mitigates the risks significantly. The table below sets out the central trade-offs noting the case for caution and the case for ambition.


Table 5: Central trade-offs

Design choice	Case for ambition	Case for caution
Revenue capture	Ownership and apportionment capture ongoing value, not one-time access; resist profit-shifting; corporate tax alone	Co-investment faces negotiating and diplomatic resistance; token-style yields depend on enforceable definitions; an

	will not deliver given standard tech structures	access fee is simpler and legally cleaner as a starting point
Creative sector	Structural displacement needs a forward-looking response; a framework calibrated to ongoing disruption is more durable than a one-time adjustment	Direct compensation has clear precedent and advocacy; broader mandates are harder to administer.
National-interest conditions	Comprehensive conditions capture full leverage and set AI-governance norms; sovereignty and transparency serve long-term interests	Extensive conditions raise diplomatic risk and may deter investment; some are hard to enforce; a simpler permit is more durable

The trade-offs suggest that the more cautious options are generally easier to negotiate and implement, while the more ambitious mechanisms offer potentially greater economic and distributional benefits but require more complex institutional and commercial arrangements. An access fee can be agreed quickly and will look adequate for as long as the upside remains hypothetical. The difficulty is that the upside is exactly what Australia would be giving away, and the gas precedent shows how hard it is to reclaim once the infrastructure is built.

Figure 6 – The Proposed Layered Model and its benefits



Note: summary diagram based on Per Capita analysis.

The layered model is designed to make ambition negotiable by sequencing it; the access fee and reserved compute can be settled early, while the ownership stakes and apportionment are negotiated as part of the investment terms, when Australia's leverage is at its peak. The greater policy risk is not that Australia negotiates a durable public interest before the value of AI infrastructure is known, but that it waits until that value is known and finds that the opportunity to negotiate such an interest has passed.

13 Conclusion

Australia has the opportunity to share in one of the largest capital programs in modern history, and it holds (some) genuine leverage over the terms. The recommendation of this report is that the country should not sell that leverage for a one-time fee. It should partner, co-own and participate ensuring Australia receives a fair apportionment of the AI dividend. Shared ownership of the fixed assets gives the public a durable claim that mobile activity and accounting profit cannot erode. Public-private partnering with real governance rights turns that claim, into influence as well as return. Revenue apportionment anchored to countable units (or a similar approach) lets the public share in commercial success without inviting the avoidance that hollowed out the PRRT. An access fee provides a floor, reserved compute provides sovereign capability, a dedicated fund preserves the value for the long term, and a properly scaled creative-sector mechanism ensures the people whose work underwrites the leverage share in the dividend. And notably this mechanism operates beyond existing copyright mechanisms, which continues to carry licensing arrangements.

Because the model is layered rather than built on any single anchor, it holds up not only across the range of capability futures (both GO-Science and Per Capita scenarios) but also against changes in how AI is built and run. If inference disperses to the edge, training stays sited and the core layers are untouched. If training itself fragments across distributed compute, capture shifts from the campus to the shared inputs the dispersed nodes still depend on, the energy, the grid and any reserved Australian capacity. A regime resting on one mechanism could be designed around and often is, a layered one that owns and conditions the inputs beneath the activity is far harder to escape. That durability is the point of building the model this way. The experience of LNG should not be repeated again.

It is recognised that these outcomes are not simple to negotiate, but it will be much harder to negotiate later, and may become implausible if we wait too long, we should not let VUCA conditions lead to indecision. The window in which Australia holds some modest leverage is open now, but it may not stay open for a very long time as there are many dynamic factors at play. The choice the Assistant Minister described, between shaping the boom and reacting to it, is in the end a choice about whether Australia designs these arrangements while it still can or is left with no dividend because of the indecision that ensued.

Appendix A: The Attributable Influence Principle

The attributable influence principle is a conceptual development that sits behind plausible future revenue apportionment and creative-sector layers (future apportionment beyond copyright) discussed in broad terms. Herein we indicate where those mechanisms might evolve as the underlying science matures. It is forward-looking rather than a recommendation for immediate implementation, and it does not alter the report's position that copyright licensing is handled through Australia's existing provisions. The principle described here concerns how contribution might in future be measured and recognised, which is a complement to licensing, not a replacement for it.

Historically, copyright has relied on the ability to identify direct acts of reproduction. A creator's legal and economic interests arise because a work has been copied, distributed, performed or otherwise reproduced in an identifiable form. Generative artificial intelligence challenges this paradigm, because the training process transforms vast collections of copyrighted material into distributed statistical representations. While these representations may not constitute recoverable copies of individual works, they nevertheless incorporate the informational content of those works into the model's learned capabilities.

The rapid development of attribution science suggests that this transformation need not sever the connection between creators and the value subsequently generated by artificial intelligence. Emerging techniques, including influence functions (Koh and Liang, 2017), gradient-tracing methods such as TracIn (Pruthi et al., 2020), data model and projection-based data attribution (Ilyas et al., 2022; Park et al., 2023), mechanistic data attribution (Wang et al., 2026), activation patching (Zhang and Nanda, 2024), and causal abstraction frameworks (Geiger et al., 2021), increasingly enable researchers to estimate the contribution of training examples to internal model representations, capabilities and, in some cases, individual outputs. We acknowledge that much of the cited literature sits adjacent to LLMs but it nonetheless offers a promising path forward, and an alternative way of thinking about knowledge construction beyond the norms of copyright.

These developments suggest that attribution should no longer be viewed as a binary concept. Rather than asking whether a copyrighted work has been directly reproduced, it becomes possible to ask whether it has made a measurable contribution to the model's productive capacity. This distinction gives rise to what is proposed here as the Attributable Influence Principle.

The Attributable Influence Principle. Where a work makes a measurable causal contribution to the productive capability or outputs of an artificial intelligence system, that contribution creates a prima facie basis for moral recognition and, where economically material, participation in the value generated by that system.

Importantly, this principle does not require perfect attribution. Modern economic systems routinely allocate value using probabilistic estimates rather than exact causal measurement. Music royalty distribution, public lending rights, collective copyright licensing, patent pools, environmental markets, and cooperative game theory all recognise that contributions may be estimated rather than observed directly. The relevant question is therefore not whether attribution can ever be perfect, but whether it is sufficiently reliable to support fair and efficient allocation.

This represents a fundamental shift in copyright theory. Copyright has traditionally been concerned with reproduction. The emergence of attribution science makes it possible to recognise contribution as an independent source of moral and economic claims. Under this framework, copyrighted works need not appear verbatim in generated outputs to retain normative significance. Instead, the contribution of those works to the learned capabilities of an AI model becomes the relevant object of recognition.

This shift matters because contemporary AI deployment is, in effect, structured to defeat the traditional test. Copyright's derivative-work doctrine depends on a traceable causal chain running from an identifiable source work to an identifiable output. The architecture of generative systems breaks that chain by design rather than by accident. Training dissolves millions of works into distributed statistical representations in which no single source persists in recoverable form; outputs are then synthesised from those diffuse representations rather than copied from any one input; and the intervening steps are opaque even to the developer. The result is that for almost any given output, no particular work can be shown to be the thing it was derived from, and the reproduction-based claim simply has nothing to attach to. Whether or not this severing is a deliberate design goal, its legal effect is the same, it places AI outputs beyond the reach of a doctrine built on identifiable copying.

The Attributable Influence Principle finds a new (plausible) avenue precisely where the old approach is impeded. It does not attempt to reconstruct the broken chain between a source work and a particular output, which is the link the architecture is designed to dissolve. It relocates the claim one level up, from the output to the capability. The question is no longer whether a given output was copied from a given work, but whether a given work made a measurable causal contribution to the model's productive capacity in the first place. That contribution survives the severing, because it is a property of how the model was built rather

than of any individual generation, and attribution science is increasingly able to estimate it. In other words, the same transformation that severs the output-level causal chain leaves the capability-level contribution intact and, with the right techniques, measurable. The claim migrates to the layer where the connection still exists.

The implications extend beyond copyright enforcement, and again for emphasis we acknowledge the limitations of extant practice but also the emerging literatures that offer promise. Attribution becomes enabling infrastructure for a broader ecosystem of AI governance. Model developers could provide transparency about the categories of works that most strongly influenced particular capabilities. Regulators could audit the provenance of specialised knowledge embedded within foundation models. Collective management organisations could administer remuneration pools based on attributable influence rather than direct reproduction. Most importantly, creators would no longer face a binary choice between uncompensated use and litigation based solely on copying. Instead, attribution science would provide a continuum along which recognition and remuneration could be calibrated according to measurable contribution.

As attribution techniques continue to improve, the distinction between explainability and governance is likely to diminish. Methods originally developed to understand neural networks may become the evidentiary foundation for new systems of copyright administration, licensing and revenue sharing. In this sense, attribution science is not merely a technical advance in machine learning. It represents the emergence of a new institutional capability that may redefine how societies recognise and reward intellectual labour in the age of generative artificial intelligence. We acknowledge that these matters are beyond the scope of the current report but must be monitored by regulators as a source of future attribution, and for discussions beyond copyright.

Relationship to the recommendations in this report

The connection to the layered model is specific and worth stating precisely, so that the appendix is not read as reopening the copyright question the body of the analysis sets aside. Two of the report's mechanisms are framed around contribution rather than reproduction, and both would become more precise as attribution science matures. The revenue apportionment in Section 5 is anchored to countable licence units today precisely because contribution cannot yet be measured cleanly enough to charge against; a mature attribution capability would, in time, allow the same apportionment to be calibrated to measurable influence rather than to a proxy. The creative-sector mechanism in Section 8, which is concerned with structural displacement and is deliberately distinct from copyright licensing, could likewise use attributable influence to direct support toward the creators and categories of work that most contributed to the capabilities generating the value, rather than distributing on a flat or purely circulation-based basis.

None of this is a precondition for acting now. The recommendations in the body of the report are designed to work with the institutional tools that exist today, and they should be negotiated while Australia’s leverage holds, regardless of how quickly attribution science develops. The Attributable Influence Principle is included here because it indicates the direction in which the apportionment and creative-sector layers can be refined over time, and because the transparency and provenance obligations that attribution would require align with the national-interest conditions recommended in Section 9. It is a horizon for the model, for discussion and future consideration, not a dependency of future negotiation or action.

References (Appendix A)

- Koh, P. W. and Liang, P. (2017), ‘Understanding Black-box Predictions via Influence Functions’, Proceedings of the 34th International Conference on Machine Learning (ICML), PMLR 70, pp. 1885–1894.
- Pruthi, G., Liu, F., Kale, S. and Sundararajan, M. (2020), ‘Estimating Training Data Influence by Tracing Gradient Descent’ (Tracln), Advances in Neural Information Processing Systems (NeurIPS) 33, pp. 19920–19930.
- Ilyas, A., Park, S. M., Engstrom, L., Leclerc, G. and Madry, A. (2022), ‘Datamodels: Predicting Predictions from Training Data’, Proceedings of the 39th International Conference on Machine Learning (ICML).
- Park, S. M., Georgiev, K., Ilyas, A., Leclerc, G. and Madry, A. (2023), ‘TRAK: Attributing Model Behavior at Scale’, Proceedings of the 40th International Conference on Machine Learning (ICML).
- Zhang, F. and Nanda, N. (2024), ‘Towards Best Practices of Activation Patching in Language Models: Metrics and Methods’, International Conference on Learning Representations (ICLR).
- Geiger, A., Lu, H., Icard, T. and Potts, C. (2021), ‘Causal Abstractions of Neural Networks’, Advances in Neural Information Processing Systems (NeurIPS) 34.
- Wang, H., Xu, Q., Wang, C., Xue, T., Peng, C., Chen, W. and Lin, F. (2026), ‘Mechanistic Data Attribution: Tracing the Training Origins of Interpretable LLM Units’, International Conference on Machine Learning (ICML).

Appendix B: Sustaining human knowledge production in the AI economy

Australia's copyright system has historically balanced two complementary objectives: rewarding creators for their intellectual labour while maintaining incentives for the continued production and dissemination of knowledge. The emergence of generative artificial intelligence presents a new policy challenge, because it affects not only how existing knowledge is used during model training, but also how future knowledge is produced (Arrow, 1962; Romer, 1990).

The current debate has focused primarily on the use of copyrighted material during the development of foundation models. While this is an important consideration, it represents only one component of a broader economic relationship between AI systems and the knowledge economy. Once deployed, generative AI becomes an active participant in markets for information, education, entertainment, professional advice and software. In each of these markets, AI-generated content competes directly with human-produced content for audience engagement.

This distinction is significant because the principal economic constraint in digital knowledge markets is no longer the production of information, but the availability of human attention. Digital technologies have progressively reduced the cost of creating and distributing content, while advances in artificial intelligence have reduced those costs further still. Human attention, however, remains fixed. Individuals have a finite capacity to read, watch, listen and learn. As Herbert Simon observed in 1971, a wealth of information creates a poverty of attention. As AI-generated content expands, it necessarily competes for a finite pool of engagement that has historically supported journalists, researchers, educators, artists, software developers, publishers and other knowledge producers.

The consequence is not simply increased competition. It is a structural redistribution of attention within the economy. AI systems are capable of generating large volumes of substitute content at very low marginal cost, capturing readership, viewership, subscriptions, advertising revenue and other forms of economic value that have traditionally sustained human knowledge production. This displacement may occur irrespective of whether individual outputs infringe copyright, because the economic effect arises through substitution rather than reproduction. This is precisely the structural displacement that the creative-sector mechanism in Section 8 is designed to address, and the reason that mechanism is framed as distinct from copyright licensing rather than as a form of it.

B.1 The circular flow of the knowledge economy

This issue warrants particular policy attention because the capabilities of generative AI are themselves derived from extensive training on human-created knowledge. The same books, articles, software, research publications, journalism, music and creative works that contribute to model capability subsequently compete against AI-generated outputs for finite audience attention. As a result, the knowledge economy experiences a circular flow in which human intellectual production strengthens AI capability, while AI deployment may reduce the future economic returns available to the creators upon whom continued knowledge production depends.

This dynamic has broader implications for productivity and innovation. High-quality research, investigative journalism, educational materials, software and creative works generate substantial public value beyond their immediate commercial returns, through knowledge spillovers and positive externalities (Arrow, 1962; Romer, 1990). If the economic incentives supporting these activities are progressively weakened through sustained displacement of audience attention, there is a risk of underinvestment in the very knowledge assets upon which future economic growth, and future generations of AI systems, depend.

For this reason, policy responses to generative AI should consider both stages of the AI value chain. The first concerns the incorporation of human intellectual property into AI capability during model development. The second concerns the downstream redistribution of economic value arising from AI's increasing share of finite human attention. Together, these effects provide a broader economic rationale for the report's value-capture and creative-sector mechanisms than copyright alone. They recognise not only the contribution of human-created works to AI capability, but also the importance of maintaining sustainable incentives for the continued production of the knowledge assets upon which Australia's future prosperity will increasingly rely.

Viewed in this way, support for the creative and knowledge sectors is not solely a question of recognising past contributions. It is also an investment in future knowledge production. A policy framework that acknowledges the role of human intellectual labour in both enabling AI capability and sustaining the knowledge economy is more likely to support continued innovation, maintain public confidence in AI adoption, and preserve the long-term foundations of Australia's research, education and creative sectors.

B.2 Information quality and market failure

The effects described above may be compounded by information asymmetries within digital content markets. As generative AI substantially increases the volume of available content, consumers may find it increasingly difficult to distinguish between authoritative, original human-created works and synthetic alternatives before consumption. Where such information

asymmetries become significant, market outcomes may increasingly resemble those described by George Akerlof in 1970 in his analysis of markets for goods of uncertain quality. In that framework, uncertainty about quality causes markets to reward average rather than superior products, reducing incentives for investment in higher-quality production.

Applied to generative AI, this suggests an additional source of market failure. If large volumes of low-cost synthetic content reduce the economic returns available to high-quality journalism, research, educational materials, software and creative works, markets may systematically underprovide knowledge assets that generate substantial public benefit. This provides a further economic rationale for interventions that extend beyond traditional copyright objectives. While copyright-based licensing appropriately governs the incorporation of existing works into AI capability, the complementary investment and revenue-sharing mechanisms recommended in this report may also be justified to sustain the continued production of high-quality human knowledge in the presence of attention displacement and quality-related market failures.

References (Appendix B)

- Akerlof, G. A. (1970), 'The Market for "Lemons": Quality Uncertainty and the Market Mechanism', *The Quarterly Journal of Economics*, Vol. 84, No. 3, pp. 488–500.
- Arrow, K. J. (1962), 'Economic Welfare and the Allocation of Resources for Invention', in R. R. Nelson (ed.), *The Rate and Direction of Inventive Activity: Economic and Social Factors*, Princeton University Press, pp. 609–625.
- Romer, P. M. (1990), 'Endogenous Technological Change', *Journal of Political Economy*, Vol. 98, No. 5, Part 2, pp. S71–S102.
- Simon, H. A. (1971), 'Designing Organizations for an Information-Rich World', in M. Greenberger (ed.), *Computers, Communications, and the Public Interest*, Johns Hopkins Press, pp. 37–72.

Notes

Sources for the figures and policy references cited in the report, in order of appearance. Several of the figures are time-sensitive, such as investment-pipeline values and fund valuations, and reflect the position as at mid-2026.

1. Global AI infrastructure investment of more than US\$650 billion in 2026 reflects 2026 capital-expenditure guidance reported for the major United States hyperscalers (Alphabet, Amazon, Meta and Microsoft, with others), as compiled in financial reporting and analysis published in early 2026. The combined figure is variously reported in the range of approximately US\$635 to 700 billion.
2. The A\$52 billion figure is the value of the fifteen data centre projects endorsed by the New South Wales Investment Delivery Authority in March 2026 for prioritised approvals support (reported as approximately A\$51.9 billion). The wider national

pipeline of announced data centre investment is larger again, with the Department of Industry, Science and Resources noting announced investments scaling beyond A\$100 billion.

3. Australian Government, National AI Plan, Department of Industry, Science and Resources, December 2025. The Plan sets the Commonwealth's strategic direction for AI and flags an intention to work with the States and Territories to establish national data centre principles.
4. Australian Government, Expectations of data centres and AI infrastructure developers, jointly announced by the Minister for Industry and Innovation, the Minister for Climate Change and Energy, and the Assistant Minister for Science, Technology and the Digital Economy, March 2026. The five Expectations cover national interest, the energy transition, water, skills and jobs (including apprenticeships and structured training), and research, innovation and local capability (including favourable compute access for start-ups, researchers and not-for-profits).
5. Inquiry into the development and regulation of data centres in New South Wales, NSW Legislative Council Public Accountability and Works Committee. The inquiry is underway with a final report due by 30 September 2026; its terms of reference include the distribution of costs and benefits, public subsidies and State facilitation, cost recovery for network and water infrastructure, and workforce development and training.
6. New South Wales Government data centre consultation paper, March 2026 (canvassing whole-of-government approaches to regulating the sector); and Australian Energy Market Commission, draft connection reforms, March 2026 (proposing technical standards requiring large electricity users, including data centres, to remain connected and respond to grid instability during disturbances).
7. Norway's Government Pension Fund Global is reported by Norges Bank Investment Management as the world's largest sovereign wealth fund, valued at more than US\$2 trillion (approximately NOK 21.3 trillion at end-2025), which is around A\$3 trillion at mid-2026 exchange rates.
8. Jobs and Skills Australia, The Clean Energy Generation: workforce needs for a net zero economy, October 2023. The study identifies electricians as the single largest area of occupational shortage for the energy transition, with an estimated additional requirement of around 32,000 electricians by 2030 and around 85,000 by 2050. The study predates the recent surge in data centre proposals.
9. Illustrative Australian commitments include Amazon Web Services' A\$20 billion data centre investment for 2025 to 2029 (announced June 2025, described as the company's largest publicly announced global technology investment) and its earlier A\$13.2 billion commitment for 2023 to 2027; and the A\$2 billion AWS 'Top Secret Cloud' partnership with the Australian Government over the decade to 2033-34, announced in 2024. Figures are drawn from Australian Government and company announcements.

10. The Hon Anthony Albanese MP, Prime Minister of Australia, address on artificial intelligence, Sydney, July 2026 (as reported). The Prime Minister announced that mandatory standards for AI and data centres will be legislated by early 2027 through a national cabinet process, strengthening the March 2026 Expectations; that data centre developers will be required to underwrite new energy generation so that no costs are passed to households or business and at least as much energy is returned to the grid as is drawn from it; that a new Office of AI will sit within the Department of the Prime Minister and Cabinet; and that creators should retain control over whether their work is used to train AI and over its price and value, with the Attorney-General facilitating consultation on copyright and artist protections. Details will be settled through the legislative and consultation processes now underway.