

9 October 2026

BSA COMMENTS ON THE NATIONAL AI STANDARDS CONSULTATION PAPER

Submitted Electronically to the Department of the Prime Minister and Cabinet

The Business Software Alliance (**BSA**)¹ welcomes the opportunity to respond to the consultation paper *Getting it right: Building AI infrastructure that works for Australia* (**Consultation Paper**).²

BSA is the leading advocate for the global software industry. BSA members create technology solutions that power other businesses, including cloud storage services, customer relationship management software, human resources management programs, identity management services, network infrastructure services, cybersecurity solutions, and collaboration systems. Our members are on the leading edge of providing AI-enabled products and services, as well as the tools used by others in the development of AI systems and applications in Australia and globally. Our members have made significant investments in Australia, and many Australian businesses and government agencies rely on their products and services.

BSA has engaged at each major stage of the development of Australia's approach to AI, most recently in our submission to the Joint Select Committee on Artificial Intelligence.³ Our comments address the proposed authorisation requirement for AI training, the allocation of obligations between entities, safety, security and transparency conditions, contributions to local capability, research and innovation, and copyright and AI training.

Summary of BSA's Recommendations

- 1. Avoid suggesting that prior authorisation is required for AI training.** Such a requirement would have no international precedent and would be at odds with the Government's objective of attracting frontier AI training and encouraging AI adoption across the economy.
- 2. Allocate obligations according to each entity's role.** Requirements on facilities should rest with data centre owners and operators. Obligations on AI developers should be

¹ BSA's members include: Adobe, Alteryx, Amadeus, Amazon Web Services, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cloudflare, Cohere, Cohesity, Dassault Systemes, Databricks, Datadog, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Trend Micro, TriNet, Veeam, Workday, Zendesk, and Zoom Communications Inc.

² Getting it right: Building AI infrastructure that works for Australia, Department of the Prime Minister and Cabinet, September 2026, <https://www.pmc.gov.au/node/6446>.

³ BSA Submission to the Joint Select Committee on Artificial Intelligence, September 2026, <https://www.bsa.org/policy-filings/australia-bsa-submission-to-the-joint-select-committee-on-ai>.

limited to matters within their role, information and control, such as how they train, test and secure their models. The Government should clearly define the AI developers and training activities within scope of the standards.

- 3. Fill specific gaps in existing frameworks rather than enact economy-wide AI legislation.** The Government should first undertake gap analyses and follow a risk-based approach. Where existing frameworks can address the risk, it should rely on non-binding guidance and practical tools informed by international frameworks.
- 4. Articulate the purpose of information sharing and keep it voluntary.** Information sharing with the Government should proceed on a voluntary basis unless voluntary arrangements prove insufficient. Any information shared should be high-level, tied to a specific purpose, and not require the disclosure of trade secrets.
- 5. Design any AI incident reporting obligation narrowly and build on existing frameworks.** This means a high severity threshold, a confirmation trigger, a single report to a single authority, recognition of reports made under comparable regimes, and protection from public disclosure.
- 6. Build local capability through AI adoption rather than conditions on AI training.** This includes adoption incentives, workforce reskilling, priority sector initiatives, procurement reform and voluntary public-private partnerships.
- 7. Align copyright policy with the objective of facilitating AI training in Australia.** The Government should keep all options open, take reference from Singapore's computational data analysis exception, and support industry-led mechanisms for rightsholders to express preferences.

Authorisation Requirement for AI Training

The Consultation Paper states that “frontier labs who are *granted authorisation* to undertake large scale AI training in Australia will be required to adhere to specific, minimum-security and safety expectations” [emphasis added].⁴

We strongly caution against framing obligations for AI developers in this way. BSA does not oppose, in principle, obligations that ensure developers build frontier AI models safely and securely, provided those obligations are commensurate with the models' risks and capabilities. However, we urge the Government to refrain from framing obligations to suggest that entities must obtain prior authorisation to conduct AI training in Australia.

To the extent the Consultation Paper envisions requiring prior authorisation for entities engaging in AI training, it would in effect establish a licensing regime for which there is no international precedent. While other jurisdictions have imposed various obligations on AI developers and deployers, including transparency requirements, risk assessments and incident reporting, no jurisdiction has required AI developers to obtain permission to undertake AI training. Requiring

⁴ Consultation Paper (2026), p. 21.

prior authorisation to train AI models is also at odds with Australia’s policy objective of encouraging frontier AI training and infrastructure investment within Australia, and more broadly encouraging the development and deployment of AI across the economy. The Consultation Paper does not provide any details about this sort of requirement, such as which entities and activities would fall within scope. Such regulatory uncertainty would further discourage AI investment.

Recommendation: The Government should avoid suggesting that prior authorisation is required for AI training, as it would be at odds with Australia’s policy objective of encouraging frontier AI training and infrastructure investment in Australia, and with its broader objective of encouraging AI development and deployment across the economy. Such a requirement would also have no international precedent.

Allocating Obligations between Entities

The Consultation Paper recognises that data centres and AI training facilities “can involve multiple entities, including facility owners, operators, tenants and AI developers”, and asks how the Government should distribute requirements according to each entity’s ability to influence the delivery, maintenance, operation and output of those facilities.⁵

We welcome this framing, which recognises the general principle that obligations should be assigned based on the role that each entity plays along the AI value chain. BSA has long supported this role-based allocation of responsibilities, which assigns each entity obligations that it can implement based on the information it holds and the decisions it controls.⁶ It also places responsibility for each risk with the entity best positioned to address it.

Accordingly, any obligations on AI developers proposed by the Government should cover only matters they determine and have visibility over, such as how a model is trained, tested, and secured. AI developers have neither the information nor the control to meaningfully address how owners and operators maintain and run data centres and AI training facilities (and vice versa, data centre operators do not control the training of frontier AI models and should not be expected to hold information about that process). These distinctions in responsibility should be made clear. The Government should also clearly define which AI developers and training activities fall within scope of the standards. An entity that fine-tunes or otherwise adapts an existing model should not be treated as a frontier AI developer solely because of those activities, and any obligations should be confined to risks arising from its own modifications.

Recommendations: AI standards should reflect the different roles of the entities involved in data centres and AI training and assign responsibilities accordingly. Requirements relating to the delivery, maintenance and operation of facilities should rest solely with facility owners and operators, with such obligations not attaching to AI developers. Obligations on AI developers should be limited to matters within their role, information and control, such as how they train, test, and secure their models, with those obligations not attaching to data centre owners and

⁵ Consultation Paper (2026), p. 9.

⁶ Effective Accountability Along the AI Value Chain: Right-Sizing AI Responsibilities for the Role, BSA, June 2026, <https://www.bsa.org/policy-filings/effective-accountability-ai-value-chain-right-sizing-ai-responsibilities-role>.

operators. The Government should clearly define the AI developers and training activities within scope of the standards.

Safety, Security, and Transparency Conditions

The Consultation Paper asks what international approaches should inform the “safety, transparency or security conditions Australia applies to developers of frontier AI models”, how those conditions can keep pace with technological change while providing legal certainty, and what information AI developers should be required to share with the Government.⁷ We offer the following observations.

International Approaches and Keeping Pace with Technological Change

On international approaches, BSA notes that there is no settled regulatory template or approach globally to promoting AI safety. Jurisdictions approach AI regulation differently depending on their wider policy objectives: some, such as the EU, have adopted economy-wide, risk-tiered legislation, while others, such as Singapore, rely on existing law supported by non-binding guidance and frameworks.

At this juncture, BSA encourages the Government to ensure that any national-level AI legislation is designed to fill specific gaps in Australia’s existing legal frameworks, and to take a risk-based approach. AI development and deployment are already subject to a range of laws, including privacy, online safety, critical infrastructure and cybersecurity regulations. This is also consistent with the Consultation Paper’s recognition that “Australia has strong existing, largely technology-neutral legal frameworks, including sector-specific guidance and standards, that can apply to AI and other emerging technologies”.⁸ The Productivity Commission’s 2025 Report on Harnessing Data and Digital Technology (**PC Report**) similarly recommended that AI-specific regulation should be considered only where “existing regulatory frameworks cannot be sufficiently adapted to handle AI-related harms”, and suggested that the Government “undertake a comprehensive set of ‘gap analyses’ to understand the risks stemming from AI and whether these risks can be dealt with under existing legal and regulatory frameworks”.⁹ As the PC Report noted, “[c]reating overarching, economy-wide regulations before identifying the gaps in existing legal frameworks risks duplication and unnecessarily high compliance burden with little benefit”.¹⁰

An overly burdensome regime will not only result in severe compliance costs but also discourage investment in and deployment of AI across the Australian economy, which is at odds with the stated goals of the Government to embrace the AI opportunity. In this regard, a risk-based approach is essential to ensuring obligations are proportionate and fit for purpose. For example, the EU AI Act is grounded in a risk-based approach that focuses most obligations on high-risk

⁷ Consultation Paper (2026), pp. 21-22.

⁸ Consultation Paper (2026), p. 21.

⁹ Harnessing Data and Digital Technology Final Report, Productivity Commission, December 2025, <https://www.pc.gov.au/inquiries-and-research/data-digital/report/>, p. 2.

¹⁰ PC Report (2025), pp. 20-21.

systems. The key principle is that obligations should be calibrated to risk, not applied uniformly regardless of context.

Singapore offers a contrasting example of how to promote AI safety and security without enacting AI-specific legislation. Singapore's approach to AI governance relies heavily on non-binding guidance and best practices that apply across all sectors of the economy. The Infocomm Media Development Authority (IMDA) provides businesses with a structured overview of important risks, best practices and case studies through its Model AI Governance Framework, which it extended to generative AI in 2024 and to agentic AI in January 2026.¹¹ Singapore complements the framework with practical testing resources such as AI Verify, an open-source testing framework and software toolkit governed by the AI Verify Foundation, which allows businesses to assess their AI systems against internationally recognised governance principles.¹² As part of a coordinated whole-of-government approach, sectoral regulators in Singapore supplement IMDA's framework and tools with domain-specific guidance. For example, the Cyber Security Agency of Singapore issued Guidelines and a Companion Guide on Securing AI Systems in 2024, followed by an addendum on securing agentic AI in 2026.¹³ Singapore's approach offers the following lessons:

- The focus is on equipping businesses to develop and deploy AI safely and responsibly, rather than on prescribing detailed obligations. Beyond providing best practices and guidance, tools such as the AI Verify toolkit give businesses the means to test their systems and to learn from how others have done so.
- Because Singapore's guidelines and frameworks are non-binding, they can be updated as the technology changes. For example, IMDA updated its Model AI Governance Framework to account for agentic AI within months of that capability becoming commercially significant. This approach has allowed Singapore to avoid amending and recalibrating binding obligations while keeping its frameworks fit for purpose. Importantly, this gives businesses the certainty to explore AI development, deployment and adoption amid rapid technological change. Binding legislation struggles to provide this level of regulatory agility and flexibility.
- Detailed domain-specific guidance can supplement a broader, overarching framework. Regulators with existing supervisory capabilities are better placed to apply AI guidance to the risks their sectors actually face and can do so without creating a second layer of obligations on top of the law those sectors already follow. However, this requires significant coordination within the Government to avoid conflicting or fragmented policymaking.

¹¹ Model AI Governance Framework for Generative AI, IMDA and AI Verify Foundation, May 2024, <https://aiverifyfoundation.sg/wp-content/uploads/2024/05/Model-AI-Governance-Framework-for-Generative-AI-May-2024-1-1.pdf>. Model AI Governance Framework for Agentic AI, IMDA, January 2026, <https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf>.

¹² AI Verify, AI Verify Foundation, <https://aiverifyfoundation.sg/what-is-ai-verify/toolkit/>.

¹³ Guidelines and Companion Guide on Securing AI Systems, Cyber Security Agency of Singapore, October 2024, <https://www.csa.gov.sg/resources/publications/guidelines-and-companion-guide-on-securing-ai-systems/>. See also: Securing Agentic AI: An Addendum to the Guidelines and Companion Guide on Securing AI Systems, Cyber Security Agency of Singapore, June 2026, <https://www.csa.gov.sg/resources/publications/addendum-on-securing-ai-systems/>.

We also encourage the Government to refer to and build on existing international frameworks, such as the OECD AI Principles, the Hiroshima AI Process, and the NIST AI Risk Management Framework,¹⁴ instead of developing Australia-specific technical requirements. Doing so would reduce regulatory fragmentation across jurisdictions and improve safety outcomes.

Recommendations: The Government should undertake gap analyses to determine whether AI-related risks can be addressed under existing legal and regulatory frameworks and ensure that any national-level AI legislation is designed to fill specific gaps identified through that process. Any obligations should follow a risk-based approach, calibrated to the risk and impact of a specific AI use case. Where existing frameworks can address the risk, the Government should recommend safety, transparency and security measures for AI training through non-binding guidance and practical tools that it can update as the technology changes. Non-binding guidance and tools can be supplemented where needed by domain-specific guidance from sectoral regulators (in coordination with the rest of Government) and should take reference from existing international frameworks.

Potential Obligation to Share Information with Government

The Consultation Paper asks what information AI developers should be required to share with the Government, and whether AI developers should provide it proactively, on request or by publishing it in a specified form.¹⁵ However, other than a general statement that minimum safety and security requirements for authorised frontier labs "will support Australia's visibility of AI risks and harms",¹⁶ the Consultation Paper does not explain why the Government would require information from AI developers or what it intends to use such information for. Without a clearly defined purpose, it is difficult to assess what information would be relevant, how much detail the Government needs, or what safeguards should apply.

BSA supports transparency and information sharing to the extent that the information shared is relevant to its audience and proportionate to the purpose for which it is sought. However, we have significant concerns with mandatory obligations to share information with the Government. Not only can such requirements be impossible to fulfil, but they can also expose confidential information and trade secrets relating to training methods that reflect significant research and development. Detailed disclosure of training data sources and curation methods may also give bad actors a roadmap for circumventing safeguards. These risks increase with the amount and granularity of the information required.

Requirements to disclose details of training data can also assume that companies possess historical records about training data that may never have been created or may not exist. AI models are developed over long periods of time. During that process, datasets are routinely cleaned, deduplicated, updated, replaced, and rebuilt as part of normal development. Training data is transformed into tokenised representations, rather than maintained as a library of original

¹⁴ Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449, amended May 2024, <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>; Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Systems, G7, October 2023, <https://www.soumu.go.jp/hiroshimaaiprocess/en/index.html>. Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, National Institute of Standards and Technology, January 2023, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

¹⁵ Consultation Paper (2026), pp. 21-22.

¹⁶ Consultation Paper (2026), p. 21.

source documents. Vendors that supplied data may no longer exist, and companies may also be required under privacy laws to minimise or delete certain datasets over time.

In the circumstances, information sharing should proceed on a voluntary basis unless there is clear evidence that voluntary arrangements have failed to provide the Government with sufficient information to manage or respond to specific risks. Frontier AI developers already have strong commercial and reputational incentives to demonstrate how they test and secure their models, and voluntary arrangements allow the Government and developers to agree on what information is useful for clearly specified purposes. A voluntary approach also avoids the difficulty of legislating an information sharing obligation before the Government knows what information it actually needs, while the technology and associated risks continue to change rapidly.

Any information shared with the Government should also be high-level, tied to a clearly specified purpose, and calibrated to the AI developer's role in the AI value chain. AI model developers are well-placed to provide information about how they develop their models, but they generally lack information about how other companies deploy their AI tools. Accordingly, AI model developers can share details of a model's capabilities and limitations, how they tested and evaluated it, the security measures they applied, and the processes they use to identify and mitigate risks. This approach also aligns with existing information sharing mechanisms, such as the recently updated Hiroshima AI Process Reporting Framework, which distinguishes between model developers, application providers, and deployers.¹⁷ Relatedly, the Government should make clear that such information sharing would *not* require the disclosure of trade secrets (e.g., source code, model weights) or sensitive information that could compromise the security of the model.

Recommendations: The Government should clearly articulate the purpose for which it seeks information from AI developers. Information sharing with the Government should proceed on a voluntary basis unless there is clear evidence that voluntary arrangements have failed to provide the Government with sufficient information to manage or respond to specific risks. Any information shared should be high-level, relevant to a specific purpose or risk, and account for the AI developer's role in the AI value chain.

AI Incident Reporting

The Consultation Paper states that authorised frontier labs “will be required to adhere to specific, minimum-security and safety expectations, such as by disclosing defined reportable AI incidents to relevant Australian authorities”.¹⁸ It also states that the Government wants to minimise regulatory burden “by leveraging existing frameworks and policy levers whenever possible, such as the SOCI Act and the FATA”.¹⁹ BSA supports this approach. Any incident reporting obligation should build on existing frameworks rather than create a parallel burdensome and compliance-heavy regime.

¹⁷ Hiroshima AI Process Reporting Framework, OECD, <https://oecd.ai/en/transparency/overview>.

¹⁸ Consultation Paper (2026), p. 21.

¹⁹ Consultation Paper (2026), p. 21.

The Consultation Paper does not yet define what makes an AI incident reportable, nor the exact scope of a reporting obligation. We encourage the Government to consider the following guiding principles:

- The definition of a reportable AI incident should be limited to specified harms that meet a high threshold of severity, such as death, serious physical injury, or major property damage. It should also rest on objective and reasonably ascertainable criteria that separate reportable incidents from minor errors and malfunctions, such as inaccurate outputs.
- The reporting obligation should apply only once the developer has confirmed that a reportable incident has occurred, not on initial detection or suspicion.
- Consistent with our comments on allocating obligations, AI developers should report only incidents they know about that relate to matters within their control, such as how a model is trained, tested and secured.
- Where an incident is already reportable under another law, a single report should satisfy the overlapping obligations. Examples within Australia include the Notifiable Data Breaches scheme under the *Privacy Act 1988* and the cyber incident reporting obligations under the SOCI Act. Relatedly, reports should go to a single designated authority.
- AI developers should be able to satisfy any incident reporting obligation by providing reports prepared under comparable regimes, such as the EU AI Act. This avoids Australia-specific requirements that add compliance burden without improving safety outcomes.
- Incident reports should be protected from public disclosure, as they will often contain commercially sensitive information and details of vulnerabilities that bad actors could still exploit before an incident is remediated. Where sharing or disclosure is necessary, the Government should share only anonymised analysis.²⁰
- The act of reporting should also not, in itself, expose the reporting entity to increased liability. Otherwise, reporting obligations may discourage the candid and timely disclosure the Government needs to understand the incident.

Recommendation: Any AI incident reporting obligation should build on existing frameworks rather than create a parallel regime. The Government should limit reportable incidents to specified harms that meet a high threshold of severity, defined by objective and reasonably ascertainable criteria. The obligation should apply only once a developer has confirmed that a reportable incident has occurred and limited to incidents within the developer's knowledge and control. A single report to a single designated authority should satisfy overlapping obligations under existing schemes, such as the Notifiable Data Breaches scheme and the SOCI Act, and AI developers should be able to satisfy any reporting obligation by providing reports prepared under comparable international regimes. Incident reports should be protected from public disclosure, and the act of reporting should not in itself expose the reporting entity to increased liability.

²⁰ Ten Principles for International Harmonization of Cyber Incident Reporting, BSA, February 2025, <https://www.bsa.org/files/policy-filings/02182025bsacyberincidreporting.pdf>

Contributions to Local Capability, Research, and Innovation

BSA shares the Government's objective of ensuring that AI investment produces lasting benefits for Australian researchers, workers, and businesses. The most effective way to achieve this is to focus initiatives on driving AI adoption across the economy, rather than to attach conditions to AI training. Broad adoption of enterprise AI across the economy should be a strategic priority. Winning with AI is not only about creating the most powerful tools, but also about enabling the best tools to make sectors across the economy more successful and globally competitive. Countries that promote secure AI adoption in the private and public sectors most effectively will see the greatest economic benefits and stimulate further innovation. Firms locate research and engineering capability, invest in local skills and form research partnerships where there is demand for what they build, a workforce able to build it, and customers ready to use it.

We encourage the Government to pursue its local capability objectives through the following:

- **Reduce the cost of AI adoption.** The Government should ensure that Australian businesses can access cutting-edge AI tools and data, including across borders.²¹ Targeted incentives for AI adoption, particularly for small and medium enterprises, can accelerate this. Australia has done so before through the Small Business Technology Investment Boost, which provided a 20 per cent bonus deduction for eligible digital technology expenditure by businesses with turnover under \$50 million until June 2023.²² The Government should also support digital literacy and reskilling for AI-augmented roles, working with industry and training providers. This would help ensure that the estimated \$150 billion in data centre investment expected by 2030 translates into productivity gains across the economy.²³
- **Focus adoption initiatives on priority sectors, with clear objectives and measurable outcomes.** Concentrating adoption efforts on specific sectors where Australia has existing strengths, such as agriculture and mining, with defined objectives for each, would give firms a concrete basis on which to invest and give the Government meaningful indicators of progress. The Government should also expand access to relevant datasets, compute resources and regulatory sandboxes that allow companies to test AI innovations safely and responsibly. Government, industry and research institutions can then work together to address sector-specific pain points and create new value for Australian businesses. For example, an initiative in mining could focus on AI-enabled exploration, predictive maintenance and autonomous operations, with progress measured through improvements in productivity and safety.
- **Lead by example through government procurement.** Australian government agencies are significant purchasers of enterprise software, and the Australian Public Service AI Plan already commits to giving all public servants access to generative AI tools and to

²¹ Global Enterprise AI Adoption Agenda, BSA, February 2026, <https://www.bsa.org/files/policy-filings/02172026bsaglobalaiadoption.pdf>.

²² Small Business Technology Investment Boost, Australian Taxation Office, <https://www.ato.gov.au/businesses-and-organisations/income-deductions-and-concessions/income-and-deductions-for-business/deductions/small-business-technology-investment-boost>.

²³ Consultation Paper (2026), p. 4.

streamlining AI procurement processes.²⁴ The Government can build on this by removing roadblocks in the procurement process. These include duplicative reviews, where different teams request the same information about the same AI tool even for low-risk uses, and contract terms that require suppliers to obtain prior written approval before using any AI system, which can capture AI features already built into standard enterprise software.²⁵ Removing these roadblocks would accelerate AI adoption across the public sector and strengthen demand for AI capability in Australia.

- **Expand access through public-private partnerships.** Where the Government wishes to expand "access to compute capacity and the capabilities needed to use it" for "Australian start-ups, innovative small businesses, researchers, research institutions, vocation education and training (VET) system stakeholders, and not-for-profit organisations",²⁶ it can do so through public-private partnerships with AI developers. The Government has already signed memoranda of understanding (**MOUs**) with leading AI developers, covering support for local researchers and start-ups, collaboration with the Australian AI Safety Institute and workforce skilling.²⁷ These agreements provide a strong foundation for mutually beneficial relationships, in which AI developers contribute resources such as compute to Australian research institutions and start-ups, and the Government supports the conditions that make continued investment in Australia attractive. Building on these voluntary arrangements, rather than converting them into conditions of authorisation or investment, would allow the scope of each partnership to reflect what the Government and each developer can most usefully offer.

Recommendations: The Government should pursue its local capability objectives by focusing initiatives on driving AI adoption across the economy, rather than by attaching conditions to AI training. This should include targeted incentives that reduce the cost of AI adoption, support for digital literacy and reskilling, and initiatives in priority sectors where Australia has existing strengths. Each initiative should have clear objectives and measurable outcomes, and be supported by access to relevant datasets, compute resources and regulatory sandboxes. The Government should also lead by example by removing roadblocks to AI adoption in government procurement, such as duplicative reviews and contract terms requiring prior approval for any use of AI. Where the Government wishes to expand access to compute for researchers, start-ups and

²⁴ Australian Public Service AI Plan 2025, Digital Transformation Agency, November 2025, <https://www.digital.gov.au/policy/ai/australian-public-service-ai-plan-2025>.

²⁵ Under the Digital Transformation Agency's AI Model Clauses, a seller that intends to use an AI system must "notify the Buyer prior to the provision of the products and services with details of the proposed AI System technology and functionality relevant to the provision of the products and services; and obtain the Buyer's prior written approval for such use" (Cl 1.1.1). While the model clauses are optional, they can become entrenched as standard terms over time. See: <https://www.hlc.com/-/media/8ceddfbb6c67472c86078098acd7e853.ashx>.

²⁶ Consultation Paper (2026), p. 22.

²⁷ Memorandum of Understanding between the Australian Government and Microsoft on Collaboration on AI opportunities, April 2026, <https://www.industry.gov.au/publications/memorandum-understanding-between-australian-government-and-microsoft-collaboration-ai-opportunities>; New agreement on AI collaboration with Anthropic, Minister for Industry and Innovation, March 2026, <https://www.minister.industry.gov.au/ministers/timayres/media-releases/new-agreement-ai-collaboration-anthropic>.

public-interest users, it should build on its existing voluntary partnerships with AI developers rather than converting them into conditions of authorisation or investment.

Copyright and AI Training

The Consultation Paper observes that “[h]ow copyright treats AI training, provides regulatory certainty for AI companies to invest in Australia and includes strong protections for creators and rights holders will be a key part of the AI standards”, and notes that the Attorney-General “will continue to lead stakeholder consultation on options to incentivise AI training while providing strong copyright protections”.²⁸

The Government cannot consider copyright policy in isolation from Australia’s broader economic objectives and AI strategy. While it is important that Australia’s copyright policy protects rightsholders, it is equally important to ensure that it does not undermine the Government’s broader ambition to facilitate frontier AI training in Australia, which can yield “national security, strategic and economic benefits”.²⁹

The Government should ensure that its copyright policy is consistent with its stated objective and strategy of facilitating frontier AI training in Australia. As things stand, a recent study of copyright regimes in seven Asia-Pacific economies found that Australia has “the least flexible framework for AI training”,³⁰ highlighting a gap between that objective and Australia’s current copyright policy. The Productivity Commission’s 2025 report on Harnessing Data and Digital Technology is instructive in this regard. The Commission found that “Australia’s copyright settings make domestic AI training less attractive than overseas locations”, explaining that “[a]t the moment, Australia’s copyright settings make it difficult for AI developers to lawfully access some training material. Other things equal, this encourages AI developers to train their models in other jurisdictions where regulatory settings are perceived to present a lower barrier”.³¹

Beyond disincentivising frontier AI training within Australia, a restrictive copyright policy also creates barriers for Australian entities that develop and deploy AI domestically. The Productivity Commission observed that “[m]ost of the economic benefits of AI in Australia will come from using AI models rather than training them”.³² Many Australian entities, from businesses to research institutions, realise those benefits by adapting existing models to local needs, including by fine-tuning them on Australian data. The Commission recognised that “[s]ome small, specialist AI models also need access to open web data for training”, and that “[t]he lack of access to open web data also affects non-commercial uses of AI”.³³ This runs counter to the Government’s ambition, set out in the National AI Plan and restated in the Consultation Paper, “to

²⁸ Consultation Paper (2026), p. 23.

²⁹ Consultation Paper (2026), p. 4.

³⁰ The Economic Importance of Flexible Copyright Regimes for AI Training in APAC, Oxford Economics, August 2026, https://www.oxfordeconomics.com/wp-content/uploads/2026/08/20260812-AI-Copyright_Report.pdf, p. 18.

³¹ PC Report (2025), pp. 33-34.

³² PC Report (2025), p. 33.

³³ PC Report (2025), p. 39.

build an AI-enabled economy that is competitive, productive and resilient, while ensuring the benefits of AI are shared across regions, industries and communities".³⁴

In respect of copyright and AI training, we encourage the Government to draw on Singapore's approach, which incentivises AI training while providing strong copyright protections. In 2021, Singapore introduced a computational data analysis exception (**CDA exception**) to permit the use of publicly available copyright material in machine learning. Crucially, the CDA exception applies to both commercial and non-commercial uses, but only where the user has lawful access to the material. For example, a user who circumvents a paywall to access material cannot rely on the exception.³⁵ Singapore's CDA exception therefore gives AI developers legal certainty without undermining copyright protections for rightsholders. It is also consistent with Singapore's broader economic objectives of encouraging AI innovation and adoption.

There are also ongoing multi-stakeholder efforts to develop effective, consensus-based technical mechanisms that allow for the expression of preferences on the collection and processing of content in AI training. Stakeholders including publishers, recording studios, film producers, AI companies, and public interest groups are actively engaged in efforts to develop a common, voluntary, industry-led standard that allows copyright holders to express a preference that their works not be included in AI training datasets. These discussions are advancing in fora including the Internet Engineering Task Force's AI Preferences Working Group, which is developing a standard vocabulary for expressing AI usage preferences. The draft vocabulary allows content owners to state their preferences on whether their content is used for AI training, as input to generative AI models, or for search.³⁶ A companion draft sets out how these preferences can be attached to content through existing tools, including robots.txt files and HTTP headers.³⁷ These mechanisms are best developed through voluntary conversations between rightsholders and the technology industries, who are collectively better positioned than the Government to develop effective, consensus-driven technical solutions.

Recommendation: The Government should ensure that Australia's copyright policy is consistent with its stated objective of facilitating frontier AI training in Australia, and with its ambition to share the benefits of AI across the economy. Ongoing consultations on AI training and copyright should keep all policy options open. The Government should take reference from Singapore's approach to copyright and AI training, which will provide both frontier AI developers and Australian entities the legal certainty they need for AI training while preserving strong protections for rightsholders. The Government should also support industry-led efforts to develop consensus-based technical mechanisms that allow rightsholders to express their preferences on AI training.

³⁴ Consultation Paper (2026), p. 3.

³⁵ Copyright Act 2021, section 244. Following public consultations in 2024, the Ministry of Law and Intellectual Property Office of Singapore published a report confirming that rightsholders can continue to use technological measures to control access to their material. See: Summary of Key Changes to Prescribed Exceptions in Part 6, Division 1 of the Copyright Regulations 2021, December 2024, https://www.mlaw.gov.sg/files/Summary_of_Key_Changes_to_Prescribed_Exceptions_in_Part_6_Division_1_of_the_Copyright_Regulations_2021.pdf.

³⁶ A Vocabulary for Expressing AI Usage Preferences, Internet Engineering Task Force, September 2026, <https://datatracker.ietf.org/doc/draft-ietf-aipref-vocab/>.

³⁷ Associating AI Usage Preferences with Content in HTTP, Internet Engineering Task Force, August 2026, <https://datatracker.ietf.org/doc/draft-ietf-aipref-attach>

Conclusion

BSA supports the Government's ambition to make Australia a destination for frontier AI investment and to ensure that investment delivers lasting benefits for Australians. We hope our comments assist in designing conditions for AI training that are workable, internationally aligned and consistent with that ambition. We welcome continued engagement with the Department of the Prime Minister and Cabinet and the Office of AI, and we would be glad to discuss any aspect of this submission.

Yours sincerely,

Tham Shen Hong
Director, Policy – APAC