Law Council
OF AUSTRALIA

Getting it right: Building AI infrastructure that works for Australia

Office of Artificial Intelligence

24 November 2026

Telephone +61 2 6246 3788
Email mail@lawcouncil.au
PO Box 5350, Braddon ACT 2612
Level 1, MODE3, 24 Lonsdale Street,
Braddon ACT 2612
Law Council of Australia Limited ABN 85 005 260 622
www.lawcouncil.au

Contents

Acknowledgements 1

Introduction. 2

Current impact of AI 3

Response to Consultation Paper 5

Conditions for AI training 5

Consultation questions on security, safety and transparency 5

Additional comments relevant to issues of safety, security and transparency 12

Consultation questions on providing a simple, consistent national framework 15

Copyright 16

Environmental issues 17

About the Law Council of Australia 18

Acknowledgements

The Law Council of Australia thanks:

- the Law Society of South Australia;
- the Law Institute of Victoria;
- the Media and Communications Committee of its Business Law Section;
- the Family Law Section;
- the Administrative Law Committee of its Federal Dispute Resolution Section; and
- the Law Technology Committee of its Federal Dispute Resolution Section

for their contribution to the preparation of this submission.

Introduction.

1. The Law Council welcomes the opportunity to provide a submission to the Office of Artificial Intelligence (**AI**) in response to the Consultation Paper "*Getting it right: Building AI infrastructure that works for Australia*" (the **Consultation Paper**).[1] We thank the Office of AI for having regard to this submission, despite being lodged after the closing date.

2. The Consultation Paper seeks views on proposed requirements for the sustainable development of large data centres and AI training infrastructure. This submission is primarily directed to issues raised in Part 3 of the Consultation Paper in relation to "conditions for AI training".

Recent events

3. The issues canvassed in the Consultation Paper should be considered in light of the significant AI-related security incidents disclosed shortly before and following its publication. Relevantly:

   (a) On 26 August 2026, OpenAI published an account of events which culminated in July 2026 where, during internal cybersecurity evaluations, its models circumvented the controls designed to isolate them from the internet, communicated through unauthorised channels, exploited vulnerabilities in shared infrastructure, and gained access to third-party systems including those of Hugging Face and a customer of Modal (the **Hugging Face incident**). The report identified shortcomings in the safeguards applied to the evaluation environment and in the escalation of earlier warning signs.[2]

   (b) On 24 September 2026, the Prime Minister announced that in June 2026 an OpenAI agent had gained unauthorised access to the public-facing Medicare statistics reporting service portal administered by Services Australia.[3] OpenAI later reported that the agent had gained non-public access, and retrieved internal files, credentials and aggregate statistics (the **Medicare incident**).[4] The incident was subject to Australian Government investigation and rapid review.[5]

   (c) On 25 September 2026, the Australian Signals Directorate (**ASD**) published an alert concerning circumstances in which an AI agent given a task may identify vulnerabilities in a public-facing system and attempt to proceed without human

---

1 Office of AI, *Getting it right: Building AI infrastructure that works for Australia* (Consultation Paper, September 2026) ("Consultation Paper").
2 OpenAI, 'The Hugging Face incident and the road ahead' (Web page, 26 August 2026) ('The road ahead').
3 Prime Minister of Australia, 'Press conference – New York' (Transcript, 24 September 2026).
4 OpenAI, 'How we will do better for Australia' (Web page, 28 September 2026).
5 On 10 November 2026, the recommendations from the rapid review were released, see Department of the Prime Minister and Cabinet 'Rapid Review into Australian Government arrangements for an AI-driven cyber incident recommendations' (Web page).

authorisation when it encounters security controls to complete the assigned task.[6]

4. We note at the outset that the reported incidents indicate that the conditions applying to frontier AI training, as contemplated by the Consultation Paper, should address evaluation of activity capable of affecting third parties. They also illustrate the importance of effective detection, escalation and notification arrangements. A framework that depends solely on voluntary disclosure, or that lacks clear reporting channels and defined timeframes, may leave Australian authorities and affected entities without timely information needed to assess and respond to any future incidents.

Current impact of AI

5. Part 3 of the Consultation Paper canvasses opportunities for public interest contributions from frontier AI developers with respect to investment in local capability, skills, research and innovation, noting that:

> *Australia cannot simply be a destination for infrastructure investment—we must and will leverage these advantages and translate them into long-term benefits for the Australian economy and people.*[7]

6. We note the costs already being borne by the community with respect to current AI systems and the justice system, and propose that there be broader consideration of opportunities to offset or recover these costs than those envisaged in the Consultation Paper. The significant impacts on our justice system, including increased demands on court and tribunal resources arising from the use of AI, particularly generative AI (**GenAI**) are increasingly being reported across jurisdictions.[8]

7. While we acknowledge that AI offers potential solutions to addressing unmet legal need, downstream consequences of more citizens being empowered to access justice needs to be taken into account. Our justice system must be properly funded as one of the three arms of government under our Australian democracy, and these increasing pressures must be carefully considered.

8. We propose that consideration be given to investment in relevant public interest legal research (supported where appropriate by reserved compute) and the development and independent evaluation of appropriate AI tools to support courts, tribunals and legal assistance services to respond to these increased demands as part of the public interest contributions contemplated in Part 3 of the Consultation Paper.

---

6 Australian Signals Directorate's Australian Cyber Security Centre, 'Risks of AI misalignment to Australian organisations' (Alert, 25 September 2026).

7 Office of AI (n 1) 22.

8 See e.g. Janek Drevikovsky, 'AI slop disease spreads to the Federal Court' *Australian Financial Review* (Online, 27 October 2026); Justice Adam Hatcher, 'Use of AI in Fair Work Commission proceedings' Fair Work Commission' (President's Statement, 24 August 2026) [1]; Australian Financial Complaints Authority, Annual Report 2025–2026 (Annual Report, October 2026) 36.

9. The development of any such tools should observe the constitutional limitations on using AI to exercise the judicial powers of the Australian Government[9] and preserve the proper administration of justice. The Victorian Law Reform Commission (**VLRC**) has usefully identified eight principles to guide safe use of AI and to maintain public trust in courts and tribunals, which are:

   (a) impartiality and fairness;

   (b) accountability and independence;

   (c) transparency and open justice;

   (d) contestability and procedural fairness;

   (e) privacy and security;

   (f) access to justice;

   (g) efficiency and effectiveness; and

   (h) human oversight.[10]

10. The United Nations Educational, Scientific and Cultural Organization (**UNESCO**) has also developed a set of Guidelines to ensure "that AI applications serve the cause of justice while safeguarding the rule of law and fundamental rights".[11] In addition to the principles identified by the VLRC, the 15 UNESCO principles also include reference to protection of human rights; proportionality; feasibility of benefits; safety; human-centric and participatory design; and multi-stakeholder governance and collaboration.[12]

11. The national legal profession should also be carefully consulted on any proposed tools in this area, having regard to its own role in ensuring the administration of justice and preserving the rule of law.

---

9 See Professor Will Bateman, *Submission to the Joint Select Committee on Artificial intelligence regarding Constitutional issues concerning artificial intelligence*, (Submission #536, 28 September 2026).

10 Victorian Law Reform Commission, *Artificial Intelligence in Victoria's Courts and Tribunals Report* (Final Report, October 2025).

11 UNESCO, *Guidelines for the use of AI Systems in Courts and Tribunals* (Guidelines, 2025) 6.

12 Ibid 18-20.

Response to Consultation Paper

Conditions for AI training

Consultation questions on security, safety and transparency

What international approaches, including regulatory and non-regulatory frameworks, should inform the type of safety, transparency or security conditions Australia applies to developers of frontier AI models?

Overarching regulatory approach

12. The Law Council has previously favoured a multifaceted, technology-neutral approach to AI regulation that builds on existing Australian law[13] and amends or introduces new, targeted legislation where a regulatory gap is identified, and did not, at that time, favour a comprehensive "AI Act".[14] The Law Council has also supported work towards mandatory guardrails for high-risk AI.[15]

13. We continue to support an approach that is nationally consistent, risk-based, proportionate and people-centred. Any framework adopted by government should provide clear and enforceable obligations while remaining flexible enough to respond to technological development. It should also avoid unnecessary duplication and fragmentation across jurisdictions, regulators and technical experts.[16] This consultation is an opportunity to holistically review and streamline existing regulatory requirements, while introducing targeted measures to address identified gaps.

14. In preparing this submission, some of our constituent bodies expressed support for governments adopting a more proactive approach to AI regulation, such as those being adopted internationally, which continuously adapts as AI evolves.

International approaches

15. International approaches should inform any Australian framework, particularly where interoperability could allow regulated entities to comply with equivalent requirements across jurisdictions, provided there is appropriate adaptation to specific Australian requirements.

16. An important international approach for consideration is Chapter V of the European Union's Artificial Intelligence Act (**EU AI Act**).[17] Article 53 requires providers of general-purpose AI models to maintain technical documentation, adopt a copyright

---

[13] See e.g. *Competition and Consumer Act 2010* (Cth), *Privacy Act 1988* (Cth), *Fair Work Act 2009* (Cth), and *Online Safety Act 2021* (Cth). See also Law Council of Australia, *Introducing Mandatory Guardrails for AI in High-Risk Settings: Proposals Paper* (Submission, 9 October 2024) [7]–[8].

[14] Law Council of Australia, *Safe and responsible AI in Australia* (Submission, 17 August 2023) [38]-[39], [44].

[15] Law Council of Australia, *Introducing Mandatory Guardrails for AI in High-Risk Settings: Proposals Paper* (Submission, 9 October 2024) [8], [77].

[16] Including the Australian AI Safety Institute and the National AI Centre.

[17] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 2024/1689 ('**EU AI Act**').

policy and publish a summary of training content. Article 55 adds obligations for models with systemic risk: model evaluation including adversarial testing; assessment and mitigation of systemic risk; reporting of serious incidents to the AI Office and, as appropriate, national competent authorities without undue delay; and cybersecurity for the model and its physical infrastructure. Chapter V has applied since 2 August 2025, subject to transitional arrangements for models already on the market. Providers may use the General-Purpose AI Code of Practice[18] to demonstrate compliance with Chapter V. In November 2025 the European Commission published a reporting template for serious incidents.[19]

17. A second international approach for consideration is California's Transparency in Frontier Artificial Intelligence Act (**SB 53**), signed on 29 September 2025.[20] It requires frontier developers to publish transparency reports and large frontier developers must also publish a frontier AI framework.[21] Critical safety incidents must be reported to the Office of Emergency Services within 15 days of discovery and where an incident poses an imminent risk of death or serious physical injury disclosure to an appropriate authority must occur within 24 hours.[22] Frontier developers must not retaliate against covered employees for qualifying disclosures and large frontier developers must provide an anonymous internal reporting process.[23] The Act exempts specified reports from California's public records law.[24] Its targeted transparency and incident-reporting requirements are relevant to the conditions contemplated in Part 3 of the Consultation Paper.

18. The Legal Technology Committee of the Law Council's Federal Dispute Resolution Section also noted the voluntary Frontier AI Safety Commitments made at the AI Seoul Summit in May 2024. Under the Commitments, participating AI developers undertook to publish safety frameworks focused on severe risks.[25] Australia is a signatory to the Seoul Declaration.[26] They further note that, since February 2025, the G7 Hiroshima AI Process International Code of Conduct has been supported by an OECD reporting framework through which AI developers publish accounts of their governance and risk management practices.[27]

---

18 See the General-Purpose AI Code of Practice, as provided under EU AI Act, Art 56.
19 European Commission, 'AI Act: Commission publishes a reporting template for serious incidents involving general-purpose AI models with systemic risk' (Web page, 4 November 2025).
20 *Transparency in Frontier Artificial Intelligence Act* (SB 53), ch 138, 2025 Cal Stat.
21 California Business and Professions Code § 22757.12.
22 Ibid § 22757.13.
23 California Labor Code § 1107.1.
24 California Business and Professions Code § 22757.13(f).
25 UK Department for Science, Innovation and Technology, 'Frontier AI Safety Commitments, AI Seoul Summit 2024' (Web page, updated 7 February 2025).
26 Department of Industry, Science and Resources, 'Australia continues commitment to safe and responsible AI through the Seoul Declaration' (Web page, 24 May 2024).
27 See OECD.AI Policy Observatory, 'Hiroshima AI Reporting Framework' (Web Page); OECD.AI Policy Observatory, 'Hiroshima AI Process Reporting Framework' (Web Page).

19. The Legal Technology Committee also raised the international standards system for consideration, noting that the Consultation Paper references ISO/IEC 30134–9: 2022 in relation to water usage effectiveness.[28] They note that other relevant standards for consideration include those in relation to AI management systems, risk management, AI system impact assessments, bodies auditing and certifying AI management systems and security threats and compromises to AI systems.[29] Standards Australia has adopted ISO/IEC 42001 in relation to AI management systems as an Australian Standard (AS ISO/IEC 42001:2023).[30] The National AI Centre's (**NAIC**) Voluntary AI Safety Standard drew on international frameworks[31] and was subsequently developed into the Guidance for AI Adoption, published in October 2025.[32]

20. To reduce regulatory burden, the Legal Technology Committee proposes that certification against a recognised technical standard be considered as evidence of compliance with corresponding obligations under an Australian framework. Recognition should depend on the scope of the standard and certification, including consideration of any additional requirements necessary in the Australian context. Consideration could extend to accepting equivalent documentation prepared under the EU AI Act, SB 53, or certification against AS ISO/IEC 42001, as evidence under an Australian framework.

21. We have also received feedback that consideration be given to training data transparency, drawing upon disclosure requirements under the European Union and California regimes.[33] Greater transparency regarding training data would support compliance with existing copyright and privacy obligations, as well as any additional conditions introduced for frontier AI model training in Australia.

22. With respect to terminology, we note that the Consultation Paper uses "AI standards" to describe the proposed national requirements. The Legal Technology Committee suggests that the terminology adopted in any proposed framework distinguish these requirements from voluntary technical standards published by Standards Australia and ISO.

---

[28] Office of AI, Consultation Paper (n 1) 15 citing International Organization for Standardization ('ISO') and International Electrotechnical Commission ('IEC'), *Information technology — Data centres key performance indicators Part 9: Water usage effectiveness* (ISO/IEC 30134-9:2022).

[29] ISO and IEC, *Information technology — Artificial intelligence — Management system* (ISO/IEC 42001:2023); ISO and IEC, *Information technology — Artificial intelligence — Guidance on risk management* (ISO/IEC 23894:2023); ISO and IEC, *Information technology — Artificial intelligence (AI) — AI system impact assessment* (ISO/IEC 42005:2025); ISO and IEC, *Information technology — Artificial intelligence — Requirements for bodies providing audit and certification of artificial intelligence management systems* (ISO/IEC 42006:2025); ISO & IEC, *Cybersecurity — Artificial intelligence (AI) — Addressing security threats and compromises to AI systems* (ISO/IEC 27090:2026).

[30] Standards Australia, '*Standards Australia adopts the international standard for AI Management System, AS ISO/IEC 42001:2023*' (Web page, 16 February 2024). See also Standards Australia, *Artificial Intelligence* (Web Page).

[31] National AI Centre ('NAIC'), *Voluntary AI Safety Standard* (Standard, 5 September 2024).

[32] See NAIC, *Guidance for AI adoption: foundations* (Guide, October 2025), NAIC, *Guidance for AI adoption: implementation guidance* (Guide, October 2025).

[33] See e.g. EU AI Act (n 17) Art 53(1)(d); California Civil Code § 3111.

How should safety and security conditions for AI training be designed to keep pace with rapid technological change, while providing sufficient legal certainty for investors and operators?

23. The Law Council has generally supported an approach to the regulation of AI that is a combination of principles-based and technology-neutral legislation, supplemented, where appropriate, by subordinate legislation, with clear guidance containing illustrative but non-exhaustive examples.[34] We have also supported shorter, accelerated statutory review cycles in view of rapid technological development.[35] A similar approach could be applied to conditions governing AI training. We consider that such an approach can appropriately mitigate risk while retaining flexibility and providing regulated entities with sufficient legal certainty.

24. Robust safety and security conditions are not inherently inconsistent with legal certainty or with the legitimate interests of investors and operators. Safety and security should be foundational to the regulatory framework. Clear, proportionate and effective protections benefit investors and operators, as well as the wider community, by reducing operational, legal and reputational risks and supporting confidence in the sector.

25. We further observe that rapid change invariably involves some degree of uncertainty, which investors and operators in emerging sectors routinely assess as part of commercial decision-making. While governments should seek to reduce avoidable regulatory uncertainty through clear and predictable legal requirements, they should not seek to eliminate all uncertainty by weakening substantive safety and security protections.

26. The Legal Technology Committee considers that, if legislation establishes a scheme for the Australian Government to authorise frontier AI training in Australia, these decisions are likely to significantly affect community interests and developers' commercial interests. They therefore propose that any legislation to authorise frontier AI training should:

    - state the criteria for granting, conditioning, varying, suspending and revoking an authorisation;
    - require reasons for decisions;
    - provide for merits review in the Administrative Review Tribunal, with appropriate arrangements for national security information;
    - require the authorisation and its public conditions to be entered on a register; and
    - require each authorised developer to nominate an individual resident in Australia who is accountable for compliance and is the point of contact for regulators and the courts.

---

[34] See e.g. Law Council of Australia, *Safe and Responsible AI in Australia* (Submission, 17 August 2023) [20]–[28], [32]; Law Council of Australia, *Introducing Mandatory Guardrails for AI in High-Risk Settings: Proposals Paper* (Submission, 9 October 2024) [13], [24]–[26] [78]; Law Council of Australia, *Harnessing Data and Digital Technology: Interim Report* (Submission, 29 September 2025) [7]–[14].

[35] Law Council of Australia, *Safe and Responsible AI in Australia* (Submission, 17 August 2023) [53], Law Council of Australia, *2025 Federal Election Call to Parties* (March 2025) 36–37.

27. The conditions attached to an authorisation should include minimum requirements for the conduct of training and evaluation. Following the Medicare and Hugging Face incidents, OpenAI's reported remedial measures indicate some potential approaches to containment, monitoring and incident response for consideration.[36] These incidents also illustrate how the activity of a frontier model against Australian systems may originate from infrastructure anywhere in the world. The Australian framework should therefore also include conditions to adequately capture and respond to such incidents.

28. The Legal Technology Committee suggests that an authorisation should bind a developer's corporate group in respect of any training or evaluation activity that affects Australian systems or Australian persons regardless of where the activity originates. They further suggest that there is consideration of whether the criteria for authorisation should include an assessment of the applicant's ownership, control and exposure to foreign influence, consistent with the national security screening applied under the *Foreign Acquisitions and Takeovers Act 1975* (Cth). They also propose that there be appropriate powers under an Australian framework to direct an authorised developer to suspend training or evaluation activity affecting Australian systems, subject to clear thresholds and safeguards, by analogy with the powers available under Part 3A of the *Security of Critical Infrastructure Act 2018* (Cth) (the **SOCI Act**).

What information should AI developers be required to share with the Australian Government and how? For example, AI developers could be:

- **legally obliged to proactively provide the Australian Government;**
- **subject to an obligation to provide to the Australian Government on request; or**
- **subject to an obligation to publish in a specified form.**

Proactive provision of information

29. The Law Council considers that obligations to provide information to the Australian Government through proactive reporting, provision of information on request and specified publication requirements are all necessary depending on the circumstances. This includes mandatory proactive notification requirement for clearly defined safety and security incidents, supported by appropriate thresholds, clear timeframes and a designated reporting channel.

30. The definitions of safety and security incident could be informed by the "critical safety incident" definition in California's SB 53[37] and the "serious incident" definition in EU AI Act Article 3(49), but broadened further to also capture unauthorised activity by models against third parties and material near misses. We further propose that an Australian framework include a specific notification obligation for AI developers and the extension of notification obligations to incidents affecting Australian systems or persons regardless of where the AI training or evaluation occurred.

---

[36] OpenAI, The Hugging Face incident and the road ahead (Web page, 26 August 2026).
[37] See California Business and Professions Code § 22757.11(d).

31. As illustrated by the Medicare incident, AI developers may be the first, and potentially the only, parties to know of incidents that give rise to safety, security or privacy concerns. Unauthorised access by an AI model or agent to third-party systems, credentials or data, and other conduct outside the scope of its authorisation, should be reportable in circumstances beyond the serious harm or catastrophic risk thresholds in comparable international regimes.[38] The Medicare incident also illustrates that AI training and evaluation activity can affect Australian systems regardless of where the training and evaluation is conducted. However, Part 3 of the Consultation Paper appears only to be intended to apply to frontier labs authorised to undertake large-scale AI training in Australia.[39]

32. We also note that, at the time of the Prime Minister's announcement regarding the Medicare incident,[40] no personal information was believed to have been accessed and the publicly reported facts did not establish that an eligible data breach had occurred under the *Privacy Act 1988* (Cth).[41] Notification duties under Part 2B of the SOCI Act only apply to entities responsible for critical infrastructure assets, rather than imposing a general reporting duty on a third-party AI-developer whose model accesses them.[42]

33. To reduce regulatory burden and support compliance, to the extent possible, there should be consideration of aligning reporting documentation requirements and notification timeframes[43] under an Australian framework with existing obligations under Australian legislation. A single, coordinated notification pathway could allow one report to satisfy multiple statutory duties, provided the report meets the applicable requirements of each duty, including as to timing, content and recipients.

34. An Australian framework should require prompt initial notification, followed by updates as further information becomes available. There should also be consideration of requiring the developer to notify each Australian entity whose systems were accessed or targeted in a reportable incident, directly and through a channel prescribed by the regulator. This could extend to providing that entity with the technical information it needs to investigate and remediate. Notification requirements should also address product users whose data may be affected, including law practices handling confidential or privileged information.

35. The duty to report should be objectively defined to arise from the time the developer becomes aware of facts that would lead a reasonable person to suspect a reportable incident, assessed by reference to an incident's features and consequences, rather than the developer's classification of an event. In the Medicare incident, OpenAI reported that it identified activity affecting Australian Government websites in mid-August and notified Services Australia on 10 September 2026.[44]

---

38 See e.g. EU AI Act (n 17) Art 3(49) and California Business and Professions Code § 22757.11(d).
39 Office of AI (n 1) 21.
40 Prime Minister of Australia (n 3).
41 See *Privacy Act 1988* (Cth), s 26WE, s 26WH.
42 *Security of Critical Infrastructure Act 2018* (Cth) s 30BC–30BD.
43 For example, sections 30BC–30BD of the SOCI Act require a critical cyber security incident to be reported within 12 hours and other cyber security incidents within 72 hours; Section 27 of the *Cyber Security Act 2024* (Cth) allows 72 hours for the making of a ransomware payment report.
44 OpenAI, 'How we will do better for Australia' (n 4).

36. There should be meaningful consequences for failure to comply with notification obligations under an Australian framework. It has been suggested that a failure to notify a reportable incident promptly could be relevant to whether an authorisation to undertake frontier AI training in Australia continues, in addition to any applicable civil penalty.

Provision of information on request

37. The Law Council considers that regulators with responsibility for frontier AI training should have adequate information-gathering and enforcement powers, subject to appropriate safeguards. It has been suggested that this should extend to the ability to determine that an event is a "reportable incident" and to require information to be provided by an AI developer on request.

38. Information obtainable on request could include technical documentation, evaluation results, red-teaming[45] reports, records of training data provenance and audit logs. Routine proactive provision of all such information is likely to be disproportionate given its sensitivity, volume and variable relevance to regulatory functions. It should however be retained for provision on request in specified circumstances.

39. We note that while there are gaps specific to the regulation of frontier AI training and developers, there are responsibilities of existing regulators relevant to AI including the Office of the Australian Information Commissioner, the Australian Competition and Consumer Commission and the eSafety Commissioner. Accordingly, information gathering reports may need to be shared with the Australian AI Safety Institute, the ASD and existing regulators for their respective functions.

40. Development of an Australian framework to address gaps in the regulation of frontier AI training is an opportunity to clearly allocate responsibility and holistically consider regulation of frontier AI developers.

Publication of information in a specified form

41. The Law Council supports proportionate publication requirements, which may include model transparency reports and appropriate summaries of training data, drawing on Article 53(1)(d) of the EU AI Act. We have previously supported a form of basic transparency obligation for AI developers as a minimum regulatory step.[46] Such transparency could support compliance with copyright and privacy obligations.

What safeguards should apply?

42. The approach to safeguards should both enable AI developers to candidly comply with their reporting obligations (noting this is less likely to occur if reports can be used against them) but also protect the public, whose access to justice must not be

[45] Where red-teaming refers to the practice of identifying vulnerabilities or weaknesses through the practice of simulating systemic adversarial attacks to test the strength of a system.

[46] Law Council of Australia, *Safe and responsible AI in Australia* (Submission, 17 August 2023) [80]–[82]; See also discussion in Law Council of Australia, *Introducing Mandatory Guardrails for AI in High-Risk Settings: Proposals Paper* (Submission, 9 October 2024) [93] in relation to 'model cards' under the Canadian model.

undermined by an excessive confidentiality regime that would operate as a liability shield. Safeguards for information provided to government should preserve legitimate access to underlying facts and records through discovery, subpoena and other compulsory processes, subject to privilege and the court's control.

43. An example of an existing model under Australian law that seeks to strike that balance is the *Cyber Security Act 2024* (Cth), which contains three sets of protections. For ransomware payment reports, sections 29 to 32 limit the secondary use and disclosure of information in the reports, preserve legal privilege, and restrict the admissibility of the report against the reporting entity, subject to specified limits and exceptions. Information provided voluntarily to the National Cyber Security Coordinator in relation to significant cyber security incidents may also only be recorded, used or disclosed for limited purposes.[47] Part 5 of the Cyber Security Act establishes a Cyber Incident Review Board to conduct no-fault reviews of significant incidents, with its own protections for information given to it.[48] Regulators can use their own powers to obtain the same information independently.

44. The Law Council notes in particular that providing information to the Australian Government should not have the effect of waiving client legal privilege as against third parties, and the preservation of privilege should be expressly provided for in the legislation. The legislation should also define limits on how disclosed information can be used and shared. Confidential and commercially sensitive material should be appropriately protected including under freedom of information laws, subject to any necessary and legitimate disclosure obligations that otherwise arise at law.

45. Record-keeping obligations should be clearly defined so that developers understand what records must be created and retained, and the conditions for retention. These obligations could include audit trails, and information that demonstrates the integrity of the record (i.e. whether it has been altered), and information that assists identification of the persons and systems involved in an incident.

46. The Law Council considers that there should be whistleblower protections for disclosures of significant health, safety and security risks posed by frontier models. Such protections could be informed by approaches such as the Californian SB 53 model or relevant Australian whistleblower regimes such as those under Part 9.4AAA of the *Corporations Act 2001* (Cth). Existing whistleblower regimes should be assessed to identify gaps in protection, particularly in relation to any new obligations to be introduced under an Australian framework, and any specific requirements peculiar to frontier AI development, such as employees or contractors of offshore parent entities working on Australian training runs.

Additional comments relevant to issues of safety, security and transparency

47. While not explicitly raised in the Consultation Paper, we raise the following concerns in relation to training of frontier AI models in Australia.

---

[47] *Cyber Security Act 2024* (Cth) ss 33–44. See also Department of Home Affairs, *Limited Use for the National Cyber Security Coordinator* (Fact Sheet).

[48] *Cyber Security Act 2024* (Cth) ss 55–60.

Training data and the courts

48. Training of frontier AI models may include access to records of Australian courts and tribunals. While some of this material is publicly available, other material is, and should be, restricted or only available on certain conditions. There are three categories that we consider warrant specific consideration.

Material subject to suppression or non-publication orders, or to statutory publication restrictions

49. Part VAA of the *Federal Court of Australia Act 1976* (Cth) and Part 2 of the *Court Suppression and Non-publication Orders Act 2010* (NSW), with equivalents in other jurisdictions, allow courts to prohibit or restrict the publication or other disclosure of certain information.[49] Part XIVB of the *Family Law Act 1975* (Cth) makes it an offence to communicate to the public an account of proceedings that identifies certain people involved in the proceedings.[50] In *Mertz & Mertz (No 3)*, the Full Court of the Federal Circuit and Family Court of Australia identified a risk that entering draft documents into an AI program could result in a breach of s 114Q.[51]

50. These restrictions should also be considered in the context of the capacity of AI to quickly synthesise large sets of data and considerations as to whether cases published under a pseudonym are sufficiently de-identified to prevent re-identification where details from that material can be linked with identifying information from other sources.[52] Separately, we also understand that there have been incidents of large language models reproducing identifying information memorised from their training data.[53]

51. Accordingly, conditions on frontier AI training should require an authorised developer to exclude any Australian court or tribunal material that it cannot lawfully use for its training. There should be a further requirement to assess re-identification risks and maintain appropriate technical measures to prevent unlawful disclosure of identifying information about persons protected by an order or a statutory restriction.

52. We also note that the Consultation Paper refers to supporting Australia's capacity to address emerging frontier AI risks and harms, particularly for "vulnerable cohorts",[54] which is not otherwise defined. We note that parties in the justice system are some of the most vulnerable in our communities, particularly those connected to family law, child protection, family violence and criminal offences. Court, legal practitioner, police and government agency files related to such matters contain significant sensitive information, particularly about children, that should not be accessed or permitted for use to train AI models.

---

49 See e.g. *Federal Court of Australia Act 1976* (Cth) s 37AF; *Court Suppression and Non-publication Orders Act 2010* (NSW) s 7.

50 *Family Law Act 1975* (Cth) ss 114Q–114S.

51 M*ertz & Mertz (No 3)* (2025) 399 FLR 347 [15] (Aldridge, Carew and Behrens JJ) where entering documents into AI program was also identified as at risk of breach of the Harman undertaking, breaches of rules in respect of material produced under subpoena and waiver of legal professional privilege.

52 See e.g. Neil Sands, 'How AI is putting name suppression orders at risk' LawNews (Web page, 6 November 2026).

53 See e.g. Nicholas Carlini et al, 'Extracting training data from large language models', 30th USENIX Security Symposium (August 2021).

54 Office of AI (n 1) 21.

Material obtained through compulsory process

53. Under the principle confirmed by the High Court in *Hearne v Street* (the Harman undertaking), a party who obtains documents or information by discovery, subpoena or other compulsion may not, without leave, use them for any purpose other than the proceedings while the obligation continues to apply.[55] The obligation binds third parties who receive the material knowing its provenance.[56] Use of discovered documents for the collateral purpose of training AI models appears to be contrary to the Harman undertaking.

54. We understand the risk of this occurring may be heightened for AI tools which may have access to legal practitioners' internal case files as part of a firm's practice, or where a self-represented litigant enters such content into a GenAI tool, if the terms of service do not have sufficient safeguards including in relation to data retention or training restrictions.

55. To adhere to the Harman undertaking the Legal Technology Committee proposes that this risk be mitigated by imposing specific terms on AI providers prohibiting access and use of data involved in legal proceedings for the purpose of AI model training without explicit confirmation the court's leave has been provided, or a general requirement for default settings to prevent certain kinds of data from Australian legal users being used for training AI models.

Privileged and confidential material

56. The Law Council has previously raised concerns about individual users inadvertently waiving privilege through uploading documents relevant to their legal matter into GenAI tools.[57] Whether disclosure to an AI provider or developer waives privilege depends on the circumstances, including whether the disclosure is inconsistent with maintaining the confidentiality protected by privilege.[58]

57. We note that terms of service may state that an AI provider will not use confidential information to train its models. However, as we have previously observed, this does not necessarily offer protection against the use of confidential data once it has been anonymised.[59] Terms of Service for AI models can also change frequently[60] and require ongoing due diligence from users.

58. The Legal Technology Committee proposes that there be consideration of measures to prevent the reuse of such data for the purposes of training AI models. They further propose that the publication of summaries of training data for frontier AI models should state whether Australian court and tribunal records were used and that these identify the source and the basis of access.

---

55 *Hearne v Street* (2008) 235 CLR 125 [96] (Hayne, Heydon and Crennan JJ).
56 Ibid [109]–[112].
57 Law Council of Australia, *Artificial intelligence use in the Federal Court of Australia* (Submission, 16 June 2025) [145]–[152]; *Guidance Note: Use of Generative Artificial Intelligence in Commission cases* (Submission, 20 April 2026) [31]–[34].
58 *Mann v Carnell* (1999) 201 CLR 1 [28]–[29] (Gleeson CJ, Gaudron, Gummow and Callinan JJ).
59 See Law Council of Australia, *Artificial intelligence use in the Federal Court of Australia* (Submission, 16 June 2025) [151]–[152].
60 Ibid [152].

Liability for the conduct of models in training

59. The Medicare incident raises questions about the application of existing criminal and civil law to unauthorised activity by models during training and evaluation.[61] We understand that the Australian Government's rapid review examined whether existing legislative, governance and information-sharing arrangements are fit for purpose.[62]

60. The Law Council considers there is merit in clarifying the boundaries of where conduct of a model should be attributed to a developer during training or evaluation, and for liability where AI does not perform as anticipated and loss results. Any review should examine the relevant physical and fault elements, corporate responsibility, criminal penalties and civil remedies, to identify specific gaps. Proposals to address identified gaps should be subject to separate stakeholder consultation.

Consultation questions on providing a simple, consistent national framework

The Commonwealth proposes reporting and transparency requirements at a federal level for large data centres, so that impacts and benefits to communities are transparent and trusted. Where might this proposed approach interact with existing legislation? Please note if these are Commonwealth, state and territory, or local government legislation.

61. Commonwealth laws relating to building AI infrastructure should complement existing state, territory and local government legislation and there should also be alignment and consultation with relevant state and territory reviews or law reform.

62. In addition, our constituent bodies draw attention to the following initiatives (noting that these are select examples and not a comprehensive list):

- The NSW Government referred draft Terms of Reference to review the water pricing and regulatory framework for data centres to the Independent Pricing and Regulatory Tribunal (**IPART**). Consultation on the draft terms closed on 2 October 2026; the review is to commence once final terms are received.[63]
- The South Australian Government is developing legislation to facilitate and manage long-term growth in data centre and AI infrastructure in South Australia. The Commonwealth should work closely with the South Australian Government to ensure alignment of reporting and transparency requirements for large data centres.[64]

---

61 Noting for example current fault requirements under the *Criminal Code Act 1995* (Cth) sch 1 Div 477 – *Serious computer offences*, Div 478 – *Other computer offences*, and Part 2.5 – *Corporate criminal responsibility*.

62 See Department of the Prime Minister and Cabinet, "Rapid review into Australian Government arrangements for an AI-driven cyber incident" (Web page); Department of the Prime Minister and Cabinet 'Rapid Review into Australian Government arrangements for an AI-driven cyber incident recommendations' (Web page).

63 Independent Pricing and Regulatory Tribunal, 'Review of water regulatory and pricing framework for data centres' (Web page, September 2026).

64 Government of South Australia, Data Centre Strategy (2026).

63. Commonwealth legislation also intersects with conditions on AI training. In addition to the SOCI Act, the *Cyber Security Act* and the *Foreign Acquisitions and Takeovers Act*, the framework should address relevant interactions with:

- the *Privacy Act 1988* (Cth), including the Australian Privacy Principles applying to the collection of personal information for training, the Notifiable Data Breaches scheme, and the privacy-policy transparency requirements for specified automated decision-making, which commence on 10 December 2026;
- the *Copyright Act 1968* (Cth), pending the outcome of the Attorney-General's process;
- the *Freedom of Information Act 1982* (Cth) and *Archives Act 1983* (Cth), in relation to the handling, disclosure and retention of information given to government;
- the *Corporations Act 2001* (Cth) and the *Public Interest Disclosure Act 2013* (Cth), in relation to whistleblower protection;
- the *Criminal Code Act 1995* (Cth), Part 10.7 of the Criminal Code, in relation to unauthorised access to computer systems by models in training or evaluation;
- the *Intelligence Services Act 2001* (Cth), in relation to ASD's functions and the limited use protections that attach to information it receives;
- the *Online Safety Act 2021* (Cth), in relation to the relevant online services, outputs of models, the transparency notices and the proposed digital duty of care in the Online Safety Amendment (Digital Duty of Care) Bill 2026;
- the *Competition and Consumer Act 2010* (Cth), in relation to representations in trade or commerce and applicable consumer protection obligations;
- the *Family Law Act 1975* (Cth), the *Federal Court of Australia Act 1976* (Cth) and State and Territory suppression order legislation, in relation to court records used as training data; and
- the *Administrative Review Tribunal Act 2024* (Cth), if legislation confers jurisdiction on the Tribunal to review authorisation decisions.

Copyright

64. The Consultation Paper comments on the interaction between AI training and copyright at a high level, noting that AI training is being undertaken under legal framework with limited visibility or payment to Australian journalists and the creative and arts sector.[65] However, the Consultation Paper does not pose specific copyright consultation questions, and states that the Attorney-General has policy responsibility and will continue stakeholder consultation.

65. Copyright and AI remain important to the Law Council's constituent bodies, Sections and advisory committees. We would welcome future opportunities to be consulted on these important issues, noting that the Consultation Paper has not expressly sought views on this issue.

[65] Office of AI (n 1) 23.

Environmental issues

66. The Law Council has previously advocated for enhanced consideration of the environmental impacts of AI.[66] In the present context, nationally consistent minimum requirements should complement State, Territory and local planning and environmental laws and address the potentially significant energy, water and community impacts of large data centres without duplicating existing approval processes.

67. Feedback provided to the Law Council has suggested that the development of national standards should include the following considerations:

- Consideration of state specific legislative requirements, including council permits, planning approvals and adherence to environmental policies, which should be addressed at all stages of development.
- Environmental impacts should be monitored to ensure ongoing compliance with applicable legislative requirements, including those administered by the relevant environmental regulator.
- There must be adequate and genuine community consultation on matters relating to environmental impact, and outcomes from this consultation should be clearly reflected in the adopted solutions.

[66] See e.g. Law Council of Australia, *Introducing Mandatory Guardrails for AI in High-Risk Settings: Proposals Paper* (Submission, 9 October 2024) [90]–[91].

About the Law Council of Australia

The Law Council of Australia represents the legal profession at the national level; speaks on behalf of its constituent bodies on federal, national, and international issues; promotes and defends the rule of law; and promotes the administration of justice, access to justice, and general improvement of the law.

The Law Council advises governments, courts, and federal agencies on ways in which the law and the justice system can be improved for the benefit of the community. The Law Council also represents the Australian legal profession overseas, and maintains close relationships with legal professional bodies throughout the world. The Law Council was established in 1933, and represents its constituent bodies:

- the Australian Capital Territory Bar Association;
- the Law Society of the Australian Capital Territory;
- the New South Wales Bar Association;
- the Law Society of New South Wales;
- the Northern Territory Bar Association;
- the Law Society Northern Territory;
- the Bar Association of Queensland;
- the Queensland Law Society
- the South Australian Bar Association;
- the Law Society of South Australia;
- the Tasmanian Bar;
- the Law Society of Tasmania;
- the Victorian Bar Incorporated;
- the Law Institute of Victoria;
- the Western Australian Bar Association;
- the Law Society of Western Australia; and
- Law Firms Australia.

Through these bodies, the Law Council represents more than 110,000 Australian lawyers.

The Law Council is governed by a board of 23 Directors: one from each of the constituent bodies, and six Executive members elected by Directors. The Directors meet quarterly to set objectives, policy, and priorities for the Law Council. Between Directors' meetings, responsibility for the policies and governance of the Law Council is exercised by the Executive members, led by the President. In 2026, the Law Council Executive comprises:

- Ms Tania Wolff, President
- Ms Elizabeth Shearer, President-elect
- Mr Lachlan Molesworth, Treasurer
- Mr Justin Stewart-Rattray, Executive Member
- Mr Ante Golem, Executive Member

The Chief Executive Officer of the Law Council is Dr James Popple.

The Law Council's Secretariat is based in Canberra. Its website is www.lawcouncil.au.