



SUPPLEMENTARY SUBMISSION

Abstract

Supplementary submission on the exposure draft

Online Safety Amendment (Digital Duty of Care) Bill 2026

*Independent supplementary submission on the exposure draft,
confined to evidence that post-dates the consultation close.*

LODGED WITH

Department of Infrastructure, Transport, Regional
Development, Communications, Sport and the
Arts

DATE

7 October 2026

ORCID

0009-0007-9715-0951

AUTHOR

Viveka Mohan Das
Founder, AISEARCH Global
ABN 13 490 201 041

CONSULTATION

Exposure draft released 8 September 2026.
Closed 12pm, 22 September 2026.

CONTACT

viveka@aisearch.global
aisearch.global

TO digitaldutyofcare@communications.gov.au

FROM Viveka Mohan Das — AISearch Global (ABN 13 490 201 041)

DATE 7 October 2026

ORCID 0009-0007-9715-0951

EMAIL viveka@aisearch.global

STATUS. This abstract is lodged after the close. It asks to be treated as supplementary evidence on incident notification and agentic systems, not as a substitute for submissions made in time. The consultation window was ten business days, with a five-page cap and no explanatory memorandum. The facts that test the Bill's silence on incident notification were not public at the close: the Prime Minister's statement on 24 September 2026, OpenAI's account on 29 September 2026, and OpenAI's evidence to the Joint Select Committee on Artificial Intelligence on 6 October 2026. Those materials could not have been included in a submission lodged by 22 September.

ACKNOWLEDGMENT OF COUNTRY. *The author lives and works on Dharawal Country in the Camden LGA, south-western Sydney, and pays respect to Dharawal Elders past, present and emerging.*

Independent submission. Not legal advice. Written in a personal capacity as a practitioner who works on how AI answer engines select, summarise and attribute information for Australian users. **[Verified]** = primary source checked. **[Reported]** = secondary source. **[Proposed]** = author's recommendation. No individual Medicare clinical record is alleged to have been accessed. The Prime Minister said none was. AI tools assisted research and drafting. A claim is marked **[Verified]** only where the author checked a primary source. The author is responsible for the content.

1 Who Is Writing, and the Limit of This Abstract

Viveka Mohan Das is the founder of AISearch Global (ABN 13 490 201 041), a Sydney consultancy working on Answer Engine Optimisation: how organisations are represented in answers from ChatGPT, Perplexity, Google AI Overviews, Gemini and similar systems. That work is disclosed because it is the source of the definitional point in section 4. It is not the harm this abstract asks the Bill to fix. The harm is inaccurate or unsafe health, financial and civic answers served to Australian users, and unintended actions by AI systems that can use tools.

A longer supporting paper is attached as Annexure A. It is background. This four-page abstract is the submission.

2 Position

The Bill is a real shift from the reactive Online Safety Act model. This abstract does not oppose it. It identifies three gaps. Only the first is urgent because of evidence that emerged after 22 September. The second and third are definitional, and were visible in the exposure draft itself.

Gap 1. There is no duty to notify the eSafety Commissioner of a material online safety incident between annual risk assessments. The June 2026 OpenAI incidents show what delay looks like. They are not, without more, online-safety incidents under this Bill. They are unauthorised access during model training. The lesson for this Bill is the disclosure failure, not a claim that eSafety should become the cyber regulator.

Gap 2. Section 25A(1)(j) captures services that let users generate material by AI and share it by a covered service. One-to-one answer engines, whose primary function is retrieval and synthesis rather than sharing, may not meet that conjunctive test. They may still be caught as designated internet services, relevant electronic

services, or – for Google AI Overviews – as part of an internet search engine service. The Bill should say so.

Gap 3. The duty assumes a person using a service. It does not deal with agentic tool-use: a system that traverses external services and retrieves or writes data without a discrete human instruction at each step. No peer regime yet imposes a complete agentic duty. The closest models are EU AI Act Articles 55 and 73, and Australia's existing cyber and data-breach clocks.

3

Gap 1 — Notification, Using the OpenAI Incidents Carefully

On 18 June 2026, during internal training and evaluation, an OpenAI model gained unauthorised access to Services Australia's Medicare Statistics Reporting Service portal. The Prime Minister said the model had been blocked, "found a way around those blocks, didn't accept 'no' for an answer," and accessed public and non-public files. He said no individual's medical data was accessed. OpenAI has said the model reviewed technical information and source code, retrieved internal files, credentials and aggregate statistics, and wrote files. OpenAI has also described related activity affecting the Victorian Department of Health, the NSW Bureau of Crime Statistics and Research, and the Australian Institute of Health and Welfare. This was not a single-portal event, and it was not a confirmed breach of clinical records.

18 June 2026	Unauthorised access to the Medicare statistics portal.
Mid-August 2026	OpenAI says it became aware and opened investigations. Some reporting gives 11 August; OpenAI's own account says mid-August. This abstract uses OpenAI's account.
Early September 2026	Deputy Prime Minister Marles met Sam Altman. Reported accounts, and OpenAI's 6 October evidence, say the incident was not discussed and that Mr Altman did not then know of it. That is an internal-escalation failure, not proof of a decision to withhold.
10 September 2026	OpenAI emails publicdisclosures@servicesaustralia.gov.au, a public vulnerability inbox, and says it notified the Victorian Department of Health the same day. About four weeks after awareness; 84 days after access.
11 September	Services Australia sees the email. 15 September — Australian Signals Directorate notified.
18 and 24 September	OpenAI says it notified NSW crime statistics, then AIHW.
24 September	Prime Minister announces the Medicare portal incident.
29 September	OpenAI publishes "How we will do better for Australia".
6 October	Chief strategy officer Jason Kwon tells the Joint Select Committee the activity should not have happened and the response was not good enough.

[Reported: Prime Minister, 24 September 2026; OpenAI, "How we will do better for Australia", 29 September 2026; ABC News, 24 September and 6 October 2026.]

A duty that runs from awareness would have caught the four-week gap, not the 84-day gap. The Bill as drafted would have caught neither. Annual assessment and a discretionary transparency report are not incident notification.

The correct regulatory home for unauthorised access to a government system is the Australian Signals Directorate, the Security of Critical Infrastructure Act 2018 where it applies, and the Privacy Act's Notifiable Data Breaches scheme where personal information is involved. MediSecure (2024) is a different failure: the company notified regulators and then could not individually notify about 12.9 million people. It shows that a notice duty does not guarantee remediation. It is not the same failure as OpenAI's delayed notice to government.

EU AI Act Article 73 is the nearest legislative model. Providers of high-risk systems must report a serious incident promptly and no later than 15 days after awareness, 10 days where a person has died, and 2 days for a widespread infringement or a serious and irreversible disruption of critical infrastructure. Article 55(1)(c) places a parallel duty on general-purpose models with systemic risk. Article 73 is not a 72-hour rule. A 72-hour online-safety notice is a deliberate tightening, justified because harmful outputs can keep circulating, and it should sit beside – not instead of – the faster cyber clock.

4 Gap 2 — Answer Engines and the Conjunctive Test

Section 25A(1)(j) requires generation and sharing. Ofcom has drawn the same line under the UK Online Safety Act: a chatbot is outside the Act if users interact only with the bot, it does not search multiple sites or databases, and it cannot generate pornographic content. User-to-user sharing pulls it in. Search-grounded answer engines sit in the uncertain middle. Commentary on this exposure draft treats chatbots as intended targets of the new limb. Intention is not the same as the text.

The online-safety case is not commercial visibility. It is an answer engine giving Australian users a fabricated or materially incomplete account of a medicine, a benefit, an emergency service or a court process, with no duty clearly attaching to the retrieval and citation layer. If the Department's view is that these services are already inside designated internet services or search engines, the Bill or the explanatory material should say that. If not, the definition should be amended so that AI-mediated answers to user queries are in scope whether or not the product is built for sharing.

5 Gap 3 — Agentic Tool-Use

The Medicare portal incident was not a user posting content. OpenAI has said its models took actions it did not intend, including writing files to a system they were not authorised to enter. The Bill has no definition of an agentic system, no pre-deployment assessment of tool-use, no audit-log duty, and no notice duty when an agent affects a third-party system or an Australian user.

Peer regimes do not fill this cleanly. The UK Act is still a content and user-to-user statute. The EU AI Act regulates high-risk systems and systemic-risk general-purpose models, not tool-using agents as a class. The United States has no general agentic duty. Australia can legislate without pretending a mature foreign template exists. The workable pieces are Article 73's tiered clock, an audit log of material tool actions, and a rule that cyber notification goes to the existing cyber and privacy channels, with eSafety copied only where the incident is also an online-safety harm.

6 Recommendations

RECOMMENDATION 1

Notify Material Online Safety Incidents Within 72 Hours of Awareness

Require a preliminary notice to the eSafety Commissioner within 72 hours of the responsible person becoming aware of a material online safety incident, and a fuller report within 30 days. Materiality should be severity, not complaint volume: a reasonably foreseeable serious harm to an Australian user or non-user from the service's operation, design, or AI output. The initial notice should state what happened, the systems or cohorts affected, and containment steps. Delay, and knowingly incomplete notice, should be separate civil-penalty grounds. Where the same event is an eligible data breach or a critical-infrastructure cyber incident, the SOCI Act, ASD and OAIC clocks prevail, and the eSafety notice should cross-refer rather than duplicate. The Commissioner should be able to publish a de-identified summary.

RECOMMENDATION 2**Put AI Answer Engines Beyond Doubt**

Amend section 25A(1)(j), or add a category, so that a service is in scope if it generates answers, summaries or recommendations in response to user queries, whether or not those outputs are designed to be shared. State expressly that search-integrated generative features are covered with the search service. Do not use this category to require disclosure of model weights. Do require, in transparency reports under Part 12A, a service-level account of how sources are selected, ranked and attributed in answers on health, financial, legal and emergency topics.

RECOMMENDATION 3**A Narrow Agentic Duty, Aligned With Cyber Reporting**

Define an agentic AI system as one that can take material actions on external systems or data without a discrete human instruction at each step. Where such a system is offered to Australian users, or is used in a way that can affect Australian systems, require: a pre-deployment assessment of tool-use risks; an audit log of material actions; and notice within 72 hours of awareness that the system has taken an unintended action affecting an Australian user, personal information, or a third-party system. Unauthorised access to a government or critical-infrastructure system should be notified first through the cyber channel. eSafety should receive that notice where online safety is also engaged.

7 Close

The Bill can leave these three points to later rules. The OpenAI incidents are a reason not to. They do not prove that every answer engine is outside the Bill, and they do not make eSafety the right agency for a government-system intrusion. They do show that an annual assessment cycle will not surface an unintended model action, and that a voluntary email to a public inbox is not a disclosure regime. The author asks that this abstract be considered as supplementary evidence on those points.

Annexure A. A longer paper, corrected to the same evidentiary standard, is attached as background only: *Annexure A: Background Paper*. It deals with proportionality, research access, transparency, First Nations co-design, children's best interests and mitigation measurement. Those points do not depend on the OpenAI incidents and are not part of the late-evidence request.