

ANNEXURE B

Annexure B: The Bill and the Wider Statute Book

Submission of Edwin Lee & Partners

Public Consultation on the Proposed AI Governance Bill

Edwin Lee & Partners, Kuala Lumpur

with a contribution from Global Law Office, Beijing, People's Republic of China (Annexure C)

Dated 31 July 2026 | Filed at policy@ai.gov.my per the Consultation Paper

[FEEDBACK] AI Governance Bill - Public Consultation - EDWIN LEE & PARTNERS - 31 JULY 2026

Part 1: AI-generated content and the Online Safety Act 2025

B1. The Online Safety Act 2025 ("Act 866") imposes duties on licensed applications service providers and licensed content applications service providers to mitigate users' exposure to harmful content (section 13), to operate reporting mechanisms (section 16), to protect child users (section 18), to make priority harmful content inaccessible (section 19), and to prepare Online Safety Plans (section 20), under the supervision of the Malaysian Communications and Multimedia Commission, with written instructions enforceable by fines of up to one million ringgit and continuing daily penalties (sections 25 and 27). Its First Schedule defines harmful content to include child sexual abuse material, content on financial fraud, obscene and indecent content, harassing or threatening content, content inciting violence or terrorism, and content inducing a child to self-harm.

B2. The relevance to the Bill is direct: generative AI is now the dominant production technology for several of those categories. Deepfake investment scams, squarely "content on financial fraud" under item 2 of the First Schedule, are the most visible AI harm affecting ordinary Malaysians today. The Government has already stated the design intent. Replying in the Dewan Rakyat on 22 June 2026, the Minister of Digital described a two-layer approach: the Bill governs AI systems across their lifecycle as a preventive layer, while content and misuse (he referred expressly to the Online Safety Act 2025, alongside the Sexual Offences Against Children Act 2017 and the Communications and Multimedia Act 1998) remain governed by existing laws (as reported in *The Star*, 22 June 2026). We agree with that division of labour, and our recommendations are directed at making it work: two layers only function as one system if the interface between them is built. As matters stand, Act 866 reaches the platform that *carries* the content but has no purchase on the Developer or Deployer of the system that *generates* it; the Bill reaches the generator and has no content-side machinery. Neither layer alone covers the pipeline; together, they do, if they are connected.

B3. Three connections would operationalise the Minister's stated approach.

- **First**, the Commission is a natural **Sectoral Lead candidate for AI-generated content harms**: it holds the content-side enforcement machinery, the licensing relationship with the platforms, and long regulatory experience of the sector; again, the Bill's own leverage-existing-frameworks logic. It has already demonstrated the system-facing half of the role in practice: in July 2026 the Commission engaged the platform X over its Grok AI tool's capacity to generate inappropriate images, resulting in system-level guardrails being implemented: an intervention directed at the generating system itself, not merely at content carried (as described by the Minister of Communications, reported in *The Vibes*, 21 July 2026). That episode proves the capability and exposes its present limit in the same breath: informal engagement under general powers depends on a cooperative counterparty and carries no statutory process, no enforceable teeth and no transparency. Designation under the Bill would convert a working practice into law.
- **Second**, the **complaint channels should be routed**. A member of the public harmed by AI-generated content can currently complain to the Commission under section 24 of Act 866 or, once the Bill is enacted, to the Central AI Authority: two doors for one grievance, with the choice determining which regulator, which powers and which remedies apply. A referral rule between the two bodies would spare complainants a jurisdictional puzzle they should never have to solve.
- **Third**, a quieter interaction deserves a sentence in the Bill's guidance: platforms will discharge their section 13 mitigation duties largely through AI moderation systems, and those systems are themselves AI Systems deployed by the platform, making decisions that affect users' expression and access, and so fall to be classified under the Bill's own risk framework. Compliance with one statute will routinely mean deploying systems regulated by the other; the guidance should acknowledge the loop rather than leave platforms to discover it.

B4. The missing joint: nothing requires synthetic content to be identifiable as synthetic. Parliament has just finished reinforcing the content layer. The Cybercrimes Bill 2026 creates an offence of transmitting or making available content generated or manipulated by means of a computer system, resembling existing persons, objects, places, entities or events, that falsely appears to a person to be authentic or truthful, with intent to facilitate an offence (clause 23); and an offence of disseminating intimate images whose definition, per the Bill's explanatory statement, expressly includes depictions generated, manipulated, altered or synthesized through artificial intelligence, machine learning or deepfake technology, including content generated using generative artificial intelligence platforms (clause 24). The Sexual Offences against Children Act 2017 has reached realistic synthetic depictions since enactment (section 4(a)(iii)–(iv)), is incorporated by reference into the Online Safety Act's First Schedule, and its 2026 Amendment Bill extends its extraterritorial reach. The content layer, in short, is AI-ready.

B5. What no layer supplies is the connective tissue all three depend on: the Act 866 platforms must **detect** harmful synthetic content (sections 13 and 19); prosecutions under clauses 23 and 24 must **prove** content was generated or manipulated; the Bill's incident channel must **identify** the generating system, and no rule anywhere requires an AI system's output to be identifiable as synthetic in the first place. We therefore recommend the Bill include a **proportionate synthetic-content transparency obligation**: disclosure to a person interacting directly with an AI system where this is not obvious; machine-readable marking of AI-generated audio, image, video and text; and clear disclosure where generated content would falsely appear authentic, with carve-outs for evidently creative or satirical work and for assistive editing. The building blocks are already Malaysian: the ADMP Guideline's duty to inform data subjects that AI is used (para 10.2.3) supplies the disclosure seed, and clause 23's own formulation supplies the definition of the content to be marked, a formulation that, we note in answer to the Questionnaire's request for international benchmarks, tracks the deepfake definition in Article 3(60) of the EU AI Act almost word for word, whose Article 50 imposes exactly this set of duties and whose recitals explain the design logic: marking obligations on the system layer are what make the content layer's detection duties workable (Recital 136). We are not aware of any provision currently in the Paper that would do more against deepfake fraud than this one. (*R19, R23*)

Part 2: Civil liability: who compensates the person harmed?

B6. Principle 3 promises "effective redress" (Paper, p. 9), and the Paper equips the Central AI Authority with directions, corrective measures and administrative penalties. What it does not say is whether a person injured by an AI system can recover from anyone. Regulatory sanction and civil compensation are different things; the Paper currently provides for the first and is silent on the second. The Minister of Digital has himself posed the underlying question publicly, asking how rules are enforced and deterrence created when the decision-maker is no longer a human or a conventional corporate entity, and whether those who build AI models owe a duty of care to build something fit for purpose (interview, *The Star*, 17 May 2026). We think those are exactly the right questions, and this Part suggests the Bill can begin answering them now, in two narrow and achievable steps, rather than leaving them wholly to future law reform.

B7. The default position, if the Bill stays silent, is the general law, and the general law is not ready. In negligence, a claimant must prove breach and causation against systems whose internal operation even their developers may be unable to fully explain; the opacity that motivates the Bill's transparency principle is the same opacity that will defeat claimants in court. The statutory alternative is Part X of the Consumer Protection Act 1999, which imposes liability without proof of negligence on producers, own-branders and importers for defective products (section 68), defines defect by reference to the safety "a person is generally entitled to expect" (section 67), and cannot be excluded by contract (section 71). But Part X was built for goods: "product" is defined as "any goods" (section 66(1)), and whether an AI system supplied as software (still more, supplied as a cloud service) is "goods" at all is at best uncertain. Even where Part X applies, two of its defences map onto AI with uncomfortable precision: the

development-risks defence (section 72(1)(d): the state of scientific and technical knowledge was not such that the defect could reasonably have been discovered) is almost purpose-built for emergent model behaviour, and the compliance defence (section 72(1)(a): the defect is attributable to compliance with a requirement imposed under written law) means that what the Bill requires of Developers and Deployers will bear on what they can plead in court.

B8. We do not suggest the Bill resolve civil liability in full: that is a substantial law-reform exercise in its own right. We recommend two narrower things. **First**, the Bill should **state its position**: whether breach of a duty under the Bill is actionable by a person who suffers harm as a result, or expressly is not. Either answer is workable; silence is not, because it will be litigated, and because the answer will be priced into every AI supply contract in Malaysia from the day the Bill takes effect, whether Developer and Deployer duties sound in damages, and whether they can be indemnified against in B2B chains, shapes procurement immediately. **Second**, NAIIO should recommend to the responsible ministry a **review of Part X's application to AI systems**, including the "goods" question and the calibration of the development-risks defence. The Bill's harm taxonomy and the CPA's defect standard should end that exercise pointing in the same direction. Finally, one note of caution on a concept the Minister has floated for the longer term, recognising a "tech personality" for autonomous systems (The Star, 17 May 2026). The Paper's own Principle 3 states the sounder rule, and the Minister's fuller remarks point the same way: responsibility must remain attributable to identifiable persons and must not be displaced onto the AI system itself. Conferring legal personality on AI systems would, at least in the near term, function mainly as a new place for responsibility to hide, a defendant with no assets and no mind to deter, and its availability as an idea makes stating the Bill's civil-liability position more urgent, not less. (R20)

Part 3: AI outputs as evidence: section 90A of the Evidence Act 1950

B9. Section 90A of the Evidence Act 1950 makes a document produced by a computer "in the course of its ordinary use" admissible as evidence of any fact stated in it (section 90A(1)); production in the course of ordinary use may be proved by a certificate from the person responsible for the computer's management (section 90A(2)); and where a certificate is given, the computer is presumed to have been "in good working order" and "operating properly in all respects" (section 90A(4)). The section applies "whether or not there was any direct or indirect human intervention" in the document's production (section 90A(5)), and by section 90C these provisions prevail over other written law.

B10. That framework was built on an assumption that made sense when it was enacted: a computer in good working order, used ordinarily, records and computes reliably, so its output is a trustworthy account of facts fed into it. Generative AI breaks the assumption's link, not its language. A large language model that produces a confident, false statement is not malfunctioning; producing plausible text, accurate or not, is its ordinary use, and it is "operating properly in all respects" while doing so. Section 90A(5)'s indifference to human intervention, drafted for automation, now reads as squarely admitting machine-generated content. The result is a presumption of reliability attaching, in terms, to a class of document for which reliability is precisely what cannot be presumed. Section 90B lets the court weigh such documents (including by considering the manner and purpose of creation, and accuracy) so the courts have tools; but a weight-stage correction to an admissibility-stage presumption is a patch, and litigants and investigators need to know where the line will be drawn.

B11. This intersects with the Bill at two specific points. **First**, the Central AI Authority's investigation function will generate reliable technical records from reported data and findings supporting corrective action: records that will end up before courts and tribunals, tendered under section 90A. **Second**, the Bill's Tier 2 documentation, logging and traceability duties will define, in practice, what records exist about what an AI system did. We recommend those duties be **designed with their evidential destination in mind**: system logs, assessment records and incident

documentation specified so they can carry weight under section 90B: unique identification of the system and version that acted (without which no chain of accountability under Principle 3 can be traced through the frequent model updates normal in AI operation), provenance, timing, and the separation of machine-generated content from records *about* machine behaviour. We further recommend NAIIO flag to the relevant authorities that the interaction between section 90A and AI-generated documents merits review; the Bill's traceability architecture will be far more valuable if the evidential status of what it produces is settled rather than presumed. (R21)

Part 4: Training data and output ownership: the Copyright Act 1987

B12. Two copyright questions sit underneath the Bill, and the Government has itself framed them. Replying to a parliamentary question in the Dewan Rakyat on 24 February 2026, the Prime Minister described the copyright concerns around AI as falling into two aspects, the material "input" used to train AI systems and the "output" those systems generate, and stated that the Copyright Act 1987 already provides a legal basis that must be effectively enforced, with MyIPO instructed to strengthen monitoring (as reported in Malay Mail, 24 February 2026). We adopt the same input/output structure, and we agree that enforcement matters. But we respectfully suggest that, on the input side in particular, the difficulty is not under-enforcement of an existing rule; it is that the Act does not yet contain a rule to enforce in either direction, and the Bill's own text touches both questions without meaning to.

B13. Training data. Training a model involves reproducing works at scale. Under the Copyright Act 1987 ("Act 332"), reproduction is a restricted act, and the exceptions in section 13(2) were not written with machine learning in view: the fair dealing exception (section 13(2)(a)) is confined by the section 13(2A) factors, the first of which weighs the commercial character of the dealing against the user, and the transient-copy exception (section 13(2)(q)) covers copies required for "viewing, listening or utilization" of a work on a network: consumption, not corpus-building. Act 332 contains no exception addressed to text and data mining or computational analysis, unlike a growing number of jurisdictions, including within ASEAN, that have enacted one. The consequence for the Bill is concrete: the Bill claims jurisdiction over AI "designed, developed, or used in Malaysia" (Paper, p. 5), yet the copyright status of the core input activity of developing AI in Malaysia is unresolved. And Principle 5 requires that data used in AI systems be "properly sourced" with attention to "provenance" (Paper, p. 10), language that quietly presupposes an answer to a question the Copyright Act has not yet been asked. A provenance duty enforced by an AI regulator, against a copyright backdrop this uncertain, risks becoming a *de facto* licensing requirement that Parliament never legislated. We draw attention to the structural parallel with Finding 2 of the core submission: the same "properly sourced" language presupposes an unresolved answer on the data protection side as well. Principle 5 is, at present, resting on two questions rather than one.

B14. Output ownership. Act 332 defines the author of a literary work as "the writer or the maker of the works" (section 3), and copyright subsists where the author is a "qualified person" at the time the work is made (section 10(1)). The Act contains no provision for works generated by a computer without a human author, and an AI system is not a person, qualified or otherwise. The likely consequence is that purely AI-generated output attracts no copyright at all, and output with mixed human and machine contribution sits on an undefined spectrum. For a Bill whose stated purpose includes stimulating innovation and giving industry "greater certainty" (Paper, p. 3), this is a material gap in the certainty on offer: Malaysian businesses building products on generative AI cannot currently say what, if anything, they own. There is a conceptual echo the drafters may find useful: Principle 3 insists that responsibility for AI must remain attributable to identifiable persons and must not be displaced onto the AI System itself; copyright law will need to decide whether *ownership* follows the same logic, but only copyright reform can decide it.

B15. Neither question can or should be answered in this Bill. Our recommendation is coordination and restraint, and we note that the coordination concern was raised on the floor of the Dewan Rakyat in the same February exchange, precisely because MyIPO sits under a different ministry from the Bill's sponsors. NAIIO should engage the ministry responsible for copyright on both questions as part of the Bill's development; indeed the National AI Action Plan 2026-2030 already commits to a reform roadmap on AI-related intellectual property protection, beginning with expert consultations and policy review, together with permissive-IP guidelines for testing. Our recommendation is that the Bill be expressly sequenced with that roadmap, and in the meantime that the Bill and its implementing instruments be drafted so that they **do not presuppose answers the Copyright Act has not given**: in particular, Principle 5's "properly sourced" language should be given content by guidance developed jointly with the copyright authorities and the Commissioner, not left for the Central AI Authority to define unilaterally. (R22)

This document is prepared by Edwin Lee & Partners. It is submitted in response to a public consultation and does not constitute legal advice to any person. Specific advice should be sought for particular circumstances.