Edwin Lee & Partners
SUPPORTING BUSINESS THROUGH LAW

WRITTEN SUBMISSION

Written Submission on the Proposed Artificial Intelligence Governance Bill

Public Consultation, National AI Office / AI Malaysia Berhad

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

A. Introduction

1. We welcome the opportunity to respond to the public consultation on the proposed AI Governance Bill, and we thank the National AI Office for consulting at the pre-drafting stage, when structural feedback can still be useful. This submission responds to the Public Consultation Paper, the Summary and the Full Questionnaire, each dated 10 July 2026 (together, "the Paper"). We have also completed the Google Form; these documents develop the reasoning behind those answers. We note that on 28 July 2026, during this consultation, the National AI Office was formally institutionalised as AI Malaysia Berhad, the national lead AI agency under the Ministry of Digital, with the Bill reported as targeted for completion by the end of 2026 (as reported, The Edge Malaysia, 28 July 2026). References in this submission to NAIO are to that office and, where the context requires, to AI Malaysia Berhad as its successor; we file at the contact point the Paper prescribes.

2. Edwin Lee & Partners is a corporate, commercial and compliance-focused law firm based in Kuala Lumpur. A significant part of our practice involves advising Malaysian and international organisations on data protection, technology and commercial matters, including compliance with the Personal Data Protection Act 2010 ("Act 709" or "the PDPA") and the instruments issued by the Personal Data Protection Commissioner ("the Commissioner") in 2024–2025. We are also, in our own operations, an AI-first firm: we use generative AI systems in our professional work, we produce AI-generated video and voice content, and we have built and deployed digital compliance tools for our clients' direct use. We therefore respond to this consultation in two capacities: as advisers, and as an organisation that would itself be regulated by the Bill as a **Deployer** of AI systems. We declare that interest at the outset, and we draw on our own deployment experience at several points below, most directly at paragraph 23 and in Annexure A. Two further matters of authorship bear on how this submission came to argue what it argues. The firm's founder analysed the Personal Data Protection Bill 2009 at the same pre-enactment stage this consultation now occupies, asking of that Bill the question every consultation of this kind must answer: whether it would confer better protection or attract more problems (LL.M research, Faculty of Law, University of Malaya, 2010). He later co-edited an international text on the field (Beyond Data Protection: Strategic Case Studies and Practical Guidance, Springer, 2013); the submission's central thesis, that this Bill should integrate with the statute that 2009 Bill became, reflects fifteen years of advising on the consequences of that Act's drafting choices. He also built and operated two technology startups before founding the practice, and the attention this submission pays to compliance cost for smaller organisations reflects that experience.

2A. Annexure C to this submission was prepared with a contribution from **Global Law Office ("GLO")** of Beijing. Established in 1984 as the first law firm founded following the implementation of China's reform and opening-up policy, GLO is today among the largest full-service law firms in the People's Republic of China, ranked for consecutive years by Chambers and Partners, The Legal 500 and Asian Legal Business, with a practice that includes advising Chinese and international technology and AI companies. We sought GLO's input for the reason developed at Finding 3 below: the developers most affected by the Bill's extraterritorial design are the least likely to respond to a Malaysian consultation, and GLO advises precisely those organisations. The division of responsibility between the two firms is stated at Annexure C, section C2: statements of foreign law and market practice are GLO's, attributed as such and gratefully acknowledged; the Malaysian analysis, and every conclusion in this submission about the operation of the Bill, are ours alone.

1. One observation frames everything that follows. Most AI systems that matter (credit scoring, hiring tools, insurance pricing, medical triage, customer profiling, government service delivery) run on personal data. For those systems, Malaysia is not writing on a blank page. Since 2024 the Commissioner has issued a set of instruments that already regulate AI-adjacent activity in operational detail: a mandatory data breach

notification regime (Circular No. 1/2025), mandatory Data Protection Officer appointments (Circular No. 2/2025), a Data Protection Impact Assessment guideline with defined triggers and a scoring method, a guideline dealing expressly with the use of AI in automated decision-making and profiling ("the ADMP Guideline"), a data-protection-by-design guideline, and a cross-border transfer guideline with a structured Transfer Impact Assessment. These are not aspirations. Organisations are operating them now.

2. The Paper's current answer to that body of law is a single phrase: Principle 5 requires that data used in AI systems be managed "in accordance with applicable laws" (Paper, p. 10). Our central submission is that this is not enough, not because the phrase is wrong, but because it leaves two regimes to collide in practice: two risk assessments of the same system, two incident reports for the same event, two vocabularies for the same actors. Annexure B shows the same pattern repeating, in different registers, with the cyber security, online safety, consumer, evidence and copyright statutes. Each collision point can be resolved now at essentially no cost. Resolving them later, through practice and dispute, will be expensive for everyone.

B. Three principal findings

1. Most of what follows is detailed and technical. Three findings are not, and we put them first because we think they are the points on which this submission is most likely to be useful to the drafters.

Finding 1: The Bill does not say whether it binds the Government, and the answer changes the architecture

1. Malaysia's two most recent statutes in this field made opposite choices, expressly and on the face of the Acts. The PDPA **does not apply** to the Federal Government or the State Governments (Act 709, section 3(1)). The Cyber Security Act 2024 **does** bind them, while preserving them from prosecution (Act 854, section 2). The Paper does not say which course the Bill will take, although it assumes public-sector application in substance: the Authority's Enablement function is described as including "capacity building functions for the Public Sector" and practical support "for public authorities" (Paper, p. 7), and the Paper's own examples of high-consequence AI include welfare determination and government service delivery (pp. 10, 11).

2. The consequence is structural rather than merely definitional, and it runs against the grain of everything else we recommend. Our submission throughout is that the Bill should build on the PDPA: Principle 5 incorporating it, the Bill's assessment recognising the DPIA, incident filing shared with the Commissioner. **Section 3(1) removes that foundation entirely where a public authority delivers services to citizens through AI.** There is no data-protection floor beneath the Bill in the public sector. The Bill therefore has to do more work there, not less: it must supply the substantive assessment, notice and oversight standards directly, because no other instrument will. The point has become more urgent during this consultation: the National AI Action Plan 2026-2030, published on 29 July 2026, commits the Government to AI-augmented public services and reimagined government-to-citizen interactions (initiatives I3 and I4), with catalytic projects that leverage government data expressly fast-tracked for immediate funding. Public-sector AI deployment is accelerating by design at the very point where the statutory floor is absent. We recognise that amending section 3(1) is beyond this consultation's scope. That is precisely why the Bill should not be drafted as though the PDPA were universally available. *(Developed at paragraphs 20–22; R5.)*

Finding 2: Principle 5 presupposes a lawful basis for training AI on personal data that Malaysian law does not clearly provide

1. Principle 5 requires that data used in AI systems be "properly sourced", with attention to "provenance" (Paper, p. 10). For the central input activity of AI development, training or fine-tuning a model on personal data, Malaysian law does not presently give a settled answer to what proper sourcing means. The uncertainty operates at three separate levels, and we think it is more useful to the drafters to set out all three than to assert a single conclusion:

2. **Whether the PDPA applies at all.** Act 709 regulates the processing of personal data "in respect of commercial transactions" (section 2(1), and the definition in section 4). Whether personal data harvested at scale from open web sources is processed in respect of a commercial transaction is, in our view, genuinely arguable and undecided.

3. **Whether the material is personal data in the relevant sense.** Data that has been effectively anonymised falls outside the Act. What counts as effective anonymisation in the context of a model that may memorise and reproduce training examples is an unresolved technical and legal question.

4. **If the Act applies to identifiable data, what the lawful basis is.** Here the position is narrow. Processing of non-sensitive personal data requires the data subject's consent (section 6(1)(a)). The six alternatives in section 6(2) are all grounds of necessity: performance of a contract with the data subject, steps towards such a contract, compliance with a legal obligation, vital interests, the administration of justice, and the exercise of functions conferred by law. **None is a legitimate-interests ground, and none is a business-improvement or product-development ground.** Act 709 contains no general publicly-available-data exception; its only public-availability concept, section 40(1)(c), is confined to *sensitive* personal data made public by the data subject's own deliberate steps. The research exemption at section 45(2)(c) is unavailable where the data is also processed for another purpose, which commercial deployment is. And repurposing data already lawfully held is separately constrained: consent is purpose-bound, section 7(2)(c)(i) requires notice before use for a new purpose, and section 8 restricts disclosure beyond the collection purpose.

5. We raise this in the Bill's interest and in industry's. The Bill claims jurisdiction over AI systems "designed, developed, or used in Malaysia" (Paper, p. 5) while the lawfulness of the core development activity is unresolved on three independent grounds. Two problems follow for the Bill's own machinery. First, Principle 5's provenance language presupposes an answer the PDPA has not given, and a provenance duty enforced by an AI regulator against that backdrop risks hardening, through administrative practice, into a position Parliament never took. Second, the Bill's fourth harm category is "contravention of any written law" (Paper, p. 11), a category which, as we say at paragraph 32, usefully brings the PDPA inside the risk framework. Read with the above, it means a substantial class of AI *development* activity is capable of registering as harm-relevant under the Bill from its first day in force, not because a system is dangerous but because an input question is unsettled. The National AI Action Plan 2026-2030 raises the stakes on both sides of this question: it builds the data supply for AI development at national scale (a Top 100 priority dataset catalogue, a national Data Exchange, a Right-To-Data channel, and implementation of the Data Sharing Act 2025 [Act 864], per initiative E9), and it separately contemplates coordination work on PDPA revisions (initiative E14). The first makes resolving the lawful-basis question more consequential; the second suggests the door to resolving it is already ajar.

6. Our recommendation is coordination and restraint rather than resolution: this gap is not the Bill's to close. NAIO should raise with the Commissioner and the responsible ministry whether Act 709 requires a lawful basis addressed to AI development; Principle 5's "properly sourced" language should be given content jointly with the Commissioner rather than defined unilaterally; and the Bill should avoid drafting that assumes the

question is settled. We express no view here on where the balance between innovation and individual control should sit; that is properly a matter for consultation on the PDPA, not a by-product of consultation on this Bill. *(Developed at paragraphs 25–28; R10.)*

Finding 3: The Bill's most consequential audience has not been asked

1. Malaysia's stated ambition is to attract AI investment. The developers actually entering this market are, to a significant degree, foreign, and they are the parties least likely to respond to a Malaysian consultation, for reasons of language, unfamiliarity with the process and limited incentive to engage at pre-drafting stage. The result is that the Bill risks being drafted without evidence of how its central definitions land on the organisations it most needs to reach.

2. Annexure C addresses that gap. We posed five questions to practitioners advising AI developers active in this region and received substantive written responses, presented there with attribution and our observations. Two findings were significant enough to promote into this document: the Authority's information demands on foreign Developers will collide with foreign secrecy and data-export law unless the framework provides a cooperation route (¶24; R7); and a Malaysian sandbox becomes a genuine regional entry pathway if its exit outcomes are recognised across ASEAN, a mechanism the EU AI Act already operates internally, which Malaysia, holding the ASEAN AI Safety Network secretariat, is uniquely placed to propose (¶44A; R24). Three further findings refined recommendations R13, R15 and R16, including a cleaner three-layer structure for the reporting duty developed jointly from the contributor's input and our own analysis. *(Annexure C.)*

3. Beyond these three, a single theme connects everything else we say: **the Bill does not arrive on empty ground.** What the Bill needs to do, and what only the Bill can do, is connect: one vocabulary for the actors, one assessment per system, one filing per incident across all engaged regimes, one transfer standard, and drafting that does not quietly presuppose reform of other statutes.

C. Recommendations

1. Our recommendations, in the order they are developed:

	Recommendation	Where
R1	Designate the Commissioner as a Sectoral Lead for AI systems that process personal data, and codify a coordination protocol with the Central AI Authority including a single-assessment and single-filing principle	¶16–18
R2	Provide an independent review or appeal mechanism for the Authority's directions, interim risk-control measures and administrative penalties	¶19
R3	State expressly how the Developer/Deployer roles map onto the controller/processor roles under Act 709, including parties added to the processing chain at run time	¶23
R4	Harmonise the personal-use exemption with section 45(1) of Act 709, confirming it covers an individual's use only: never the provider, never a public authority	¶22
R5	State expressly whether the Bill binds the Federal and State Governments, following section 2 of Act 854; adopt a single definitional set across the public and private sectors; and narrow the national security exemption so that dual-use systems fall inside	¶20–22
R6		¶23

	Recommendation	Where
	Anchor the definition of AI to an internationally interoperable formulation, keeping the Paper's functional list as illustrations	
R7	Give the extraterritorial scope an enforcement anchor (a local representative for foreign Developers of high-risk systems, and/or the Malaysian-established Deployer as express compliance interface) and pair the Authority's information-gathering powers with a stated mechanism for foreign legal prohibitions, routing conflicts through regulatory cooperation rather than penalising legally-compelled non-production	¶24 · Annexure C
R8	Define the further terms on which the Bill's operative duties depend: AI harm, AI incident, the high-risk threshold, automated decision-making and profiling, synthetic content, substantial modification, general-purpose AI	¶23
R9	Redraft Principle 5, or provide in its first implementing instrument, to incorporate the PDPA regime expressly	¶26
R10	Do not draft Principle 5 as though a lawful basis for training AI on personal data already exists; raise the gap with the Commissioner and the responsible ministry	¶25, 28
R11	Add a sixth principle on fairness and non-discrimination, grounded in Malaysia's own regulatory materials	¶29
R12	Give the "due regard" duty a demonstrable floor for high-risk systems; define human oversight by capability and by degree of involvement rather than by presence; and recognise certification against an established AI management standard as evidence of discharge	¶30–31
R13	Expand the harm categories to include significant financial harm, unjustified discriminatory treatment, and significant reputational harm or denial of access to essential services; and qualify "contravention of any written law" by reference to gravity	¶33–34
R14	Define Tier 1 by use-case severity rather than by intent, with intent as an aggravating factor	¶35
R15	Provide a primacy rule where the three classification axes conflict: the highest applicable classification governs, with Developer and Deployer duties cumulative	¶36
R16	Adopt mutual recognition between the Bill's high-risk assessment and the PDPA DPIA; single-window incident filing across the Bill, the PDPA and Act 854 with compatible timelines and a knowledge-based trigger; specify report content as a superset of Circular No. 1/2025; structure the duty in three layers (mandatory for incidents, voluntary and safe-harboured for near misses, a mitigation duty for mere weaknesses) and designate the Deployer as primary reporting party, with responsibility allocation resolved after filing	¶37–42 · Annexure C
R17	Codify how the PDPA applies inside the AI Sandbox; give exit a defined regulatory effect; and distinguish the Paper's technology-sandbox functions from its regulatory-sandbox function	¶43–45
R18	Confirm that offshore AI data flows are governed by the existing cross-border transfer regime under Act 709	¶46
R19	Coordinate the Bill with the Online Safety Act 2025 on AI-generated harmful content; consider MCMC as Sectoral Lead for that domain; route the public-complaint channels	Annexure B
R20	State the Bill's position on civil liability, and initiate a review of whether Part X of the Consumer Protection Act 1999 reaches AI systems	Annexure B
R21	Review the interaction between AI-generated documents and section 90A of the Evidence Act 1950, and design the Bill's documentation duties so the records can bear evidential weight	Annexure B
R22	Coordinate the Bill's data-provenance expectations with copyright reform, so the Bill does not presuppose answers the Copyright Act 1987 has not given	Annexure B

	Recommendation	Where
R23	Include a proportionate synthetic-content transparency obligation as the connective tissue between the Bill, the Online Safety Act's platform duties and prosecutions under the Cybercrimes Bill 2026	Annexure B
R24	Frame the AI Sandbox additionally as an inbound-investment instrument and regional entry pathway: SME measures (free and priority access, outreach, pre-deployment guidance), exit reports carrying defined weight in subsequent classification and assessment, and a Malaysian initiative for ASEAN-level mutual recognition of sandbox outcomes, on the model the EU AI Act already operates internally	¶44A · Annexure C

D. Focus Area 1: AI Governance Architecture

1. We support central institutional oversight. The Paper's diagnosis is accurate: an organisation deploying AI across its finance, HR and customer functions can already face several regulators with different expectations, and a national baseline reduces that fragmentation rather than adding to it, provided the new Authority coordinates with the regulators that already exist.

2. The Paper anticipates this through Sectoral Leads: public bodies appointed and delegated powers where they have "sufficient legal authority, technical expertise, and governance capacity" (Paper, p. 8). We would make the obvious candidate explicit. **The Commissioner satisfies each criterion for AI systems that process personal data.** The legal authority exists under Act 709, including the production powers under section 121 which Circular No. 1/2025 itself invokes. The technical and governance capacity is demonstrated by the instruments issued since 2024, including the ADMP Guideline, which is, to our knowledge, the only Malaysian regulatory instrument that currently addresses the use of AI in decision-making about individuals in operational terms. This is not a request for something new; it applies the Bill's own design of leveraging existing frameworks (Paper, p. 8) to its most developed instance.

3. There is a recent regional demonstration that a regulator of this kind can do the work. On 20 July 2026, ten days after this consultation opened, Singapore's Personal Data Protection Commission issued advisory guidelines on the use of personal data in generative AI, addressing the lawful basis for training, when notification to users must be AI-specific rather than general, the allocation of responsibility across the model supply chain, and how access and correction rights operate against a trained model. Singapore has no comprehensive AI statute; its data protection regulator nonetheless produced operational AI governance guidance under a consent-based Act comparable to Act 709. We note that our Commissioner is in one respect further along: the ADMP Guideline addressed AI in decision-making about individuals before either jurisdiction had an AI statute in contemplation.

4. Whether or not the appointment is made, we recommend the Bill codify a **coordination protocol** between the Authority and the Commissioner rather than leaving coordination to administrative practice. The protocol needs two things: a primacy rule for genuine conflicts, and a principle that one AI system is assessed once and one event reported once, with the authorities sharing the output. If it is not in the statute, every regulated organisation will negotiate it case by case.

18A. A development during this consultation strengthens the case for codification and adds one question the Bill should answer. On 29 July 2026 the Ministry of Digital published the National AI Action Plan 2026-2030 ("the Action Plan"). Two of its governance initiatives commit to precisely the machinery we recommend: initiative E11 lists, among its implementation priorities, the formalisation of coordination mechanisms between the central AI

governance authority and sector regulators, and the enactment of a principles-based framework that codifies core definitions and risk tiers and aligns with international standards to support interoperability and investor confidence; initiative E12 commits to formal coordination protocols with sector regulators and relevant agencies covering incident reporting and AI sandboxes, to mandated visibility of AI incidents across enforcement agencies, and to a public AI incident registry. R1 and R16 ask, in substance, that these commitments be placed in the Bill rather than left to administrative arrangement; the Bill is the natural and perhaps only durable home for them. One point of architecture now needs express clarification across three documents published within nineteen days of each other. The Paper vests AI Safety, enforcement and enablement functions in a proposed **Central AI Authority** (Paper, pp. 7-8). The Action Plan establishes an **AI Trust Function** (E12) described as coordinating with that proposed authority, holding standards, evaluation and certification, incident monitoring and forensics, and international engagement functions. And at the Action Plan's launch on 28 July 2026, the Government announced that NAIO is institutionalised as **AI Malaysia Berhad**, whose mandate is reported to include developing governance frameworks and ethical standards, and under which a **Malaysian AI Safety Institute** will be established to conduct safety assessments, model testing and red-teaming, functions that appear to give institutional form to the Trust Function (as reported, The Edge Malaysia, 28 July 2026). We make three observations, each answerable in the Bill. First, the Bill should state which body is the Central AI Authority and how AI Malaysia Berhad relates to it. If the design is that AI Malaysia Berhad develops and promotes while a statutory Authority regulates, we would commend it: separating promotion from regulation avoids the structural conflict of a promoter-regulator, and Malaysian precedent supports the split (MDEC develops the digital economy; MCMC regulates it). Second, and as a matter of legal form: a Berhad is a company, and the Bill's coercive powers, directions, interim risk-control measures, administrative penalties and investigation, should be vested in a statutory authority, not in a company, however owned. Conferring enforcement powers on a corporate body raises accountability and review questions the Bill does not need to create. Third, the AI Safety Institute's assessments should be given defined status under the Bill: an Institute evaluation of a system ought to count as evidence in the Bill's risk classification and due-regard framework, on the same recognition logic we recommend for certification at R12, so that the Government's own testing capability and the Bill's compliance architecture reinforce each other. Whichever way these choices fall, the coordination protocol at R1 and the single-window principle at R16 become more important with each institutional layer, not less.

1. Separately, we invite attention to the Authority's enforcement design. The Paper gives it rule-making power, investigation power, and the power to issue directions, impose interim risk-control measures and levy administrative penalties (Paper, pp. 7–8). It says nothing about review or appeal. We do not suggest the powers are inappropriate, an AI regulator plainly needs teeth, but the combination of rule-maker, investigator and penaliser in one body, with no stated route of challenge, is unusual by Malaysian regulatory standards and will invite challenge on administrative-law grounds if left unaddressed.

19A. Finally on the architecture, one development since the Paper was issued deserves the drafters' attention, because it converts a line of the Paper into a live institutional need. The Authority's AI Safety functions include "engaging in international technical cooperation" (Paper, p. 7). Six days after the Paper's date, on 16 July 2026, Malaysia became one of 29 founding members of the World Artificial Intelligence Cooperation Organization established in Shanghai, adding a treaty-level AI cooperation commitment to the ASEAN AI safety coordination role Malaysia already holds (as reported by Bernama, 17 July 2026). We make two observations, both structural. First, this vindicates the Bill's own design: a principle-based statute with detail reserved to subsidiary instruments is precisely the architecture that can absorb standards emerging from multilateral processes (whether from that organisation, from ASEAN, or from ISO) without amending primary legislation, and the Bill should be drafted so that the Authority's cooperation function expressly extends to representing Malaysia in such bodies and to giving domestic effect, through its instruments, to standards Malaysia adopts in them. Second, Malaysia's commitments now span frameworks anchored in different regulatory traditions, and its businesses serve users subject to others. That makes definitional interoperability a practical necessity rather than a courtesy, and it strengthens the

recommendation at R6: anchoring the Bill's definitions to neutral international formulations is the only choice that serves all of Malaysia's alignments at once.

E. Focus Area 2: Scope of the Bill

1. **Does the Bill bind the Government?** For the reasons at paragraphs 6–7 above, we regard this as the most consequential question the Paper does not reach, and we do not think it can safely be left to implication.

2. Three consequences follow. First, the absence of a PDPA floor in the public sector means the Bill must supply the substantive standard directly for public-sector AI processing citizens' data. Second, and answering the Questionnaire's question directly, **the five definitions should be common across the public and private sectors.** A citizen refused a benefit by an automated system and a customer refused a loan by one are exposed to the same category of harm; a single vocabulary is what allows a single body of guidance, one assessment methodology and one incident taxonomy to serve both. Divergent definitions would produce two AI regimes in one country, with the weaker applying where the individual has least choice about participating. Third, the national security exemption should be narrowed. The Paper exempts systems used "solely for national defence or security" (p. 5); "solely" is carrying a great deal of weight, and dual-use is the norm in this field. We recommend the exemption attach to the *use* rather than the system, so that a system deployed partly outside the defence and security domain is regulated in respect of that other use.

3. **The personal-use exemption should be harmonised with section 45(1) of Act 709 deliberately, not accidentally.** The wording is nearly identical, and we assume that is intentional. Two clarifications close the gap the current drafting leaves. The exemption should attach to the individual's *use* only: the developer and deployer of a consumer AI product remain fully regulated even though every one of their users is exempt; section 45(1) works exactly this way. And a public authority delivering services through an AI system should never fall within it, including where the citizen reaches the service through a personal AI agent. As drafted, an interaction between an exempt personal agent and a government system could be argued to sit outside the Bill at both ends.

4. **Three definitional points.** *AI Systems:* the definition's closing limb, "another digital system capable of generating outputs such as predictions, recommendations, content or decisions" (Paper, p. 4), would, read alone, capture a rules-based spreadsheet model, which the Paper's own distinction from "ordinary software... based on fixed instructions" plainly intends to exclude. We can offer a live example from our own practice, declared at paragraph 2. We have built and deployed, for our clients' direct use, PDPA-readiness and DPIA questionnaire tools. They are entirely deterministic: fixed logic, decision trees and scoring rules written by our lawyers, calling no model at any point, producing identical output on identical input, yet they plainly "generate recommendations", and the closing limb as drafted arguably captures them. We built those tools *with* AI coding assistance, which illustrates the same boundary from the other side: software made with AI is not thereby software containing AI. No one, we think, intends a law firm's rules-based compliance questionnaire to be regulated as an AI System, and a short qualifier tying the closing limb back to the artificial-intelligence capability described earlier in the definition resolves it. *The AI Lifecycle:* we agree its value lies in locating the stage at which a duty bites, and that value increases if the stages are mapped against duties Malaysian organisations already owe: design and development to the data-protection-by-design guideline, the point before deployment to the DPIA, operation and monitoring to security and breach notification, retirement to retention and destruction. *Developer and Deployer:* these track controller and processor closely but not perfectly. A Deployer running an AI system on personal data will usually be a controller; a Developer providing a hosted model may be a processor for its customers' data and a controller for its own training data. "Usually" is not a drafting standard, and the mapping should confirm that one

organisation can hold both roles at once, including where a party joins the processing chain at run time, which is how systems that call external tools and services now assemble.

> **Further definitions the Bill will need (Scope Q5).** Each is a term on which an operative duty already silently depends, and for most a Malaysian instrument supplies the content: **AI harm** (tie to the harm categories, so the reporting duty and the risk framework describe the same thing); **AI incident** (currently a working description, with near misses as a distinct category); **the high-risk threshold** (Tier 2 is described by contrast, not defined; the test of a workable threshold is whether a compliance officer can apply it without external advice); **automated decision-making** and **profiling** (adopt the ADMP Guideline's content rather than restating it); **personal data** (adopt section 4 of Act 709 by reference; a second Malaysian definition would be a serious and avoidable problem); **synthetic content** (see Annexure B); **substantial modification** (without a threshold, a classification made once is treated as permanent and the assessment obligation expires on the first model update); and **general-purpose AI** (the Paper's definitions assume an identifiable intended purpose, which is the assumption a general-purpose model does not satisfy: the same model may be benign in one deployment and high-risk in another, and the Developer cannot classify by use because it does not choose the use). Annexure C addresses this last point with worked configurations.

1. **Extraterritorial scope needs an enforcement anchor, and its information demands need a conflicts mechanism.** The reach is right; what the Paper does not answer is how it will be made real against a foreign Developer with no Malaysian establishment. Directions, interim measures and penalties bind only those the Authority can reach. Malaysian law has solved this twice in the statutes closest to this one. The PDPA applies to persons not established in Malaysia who use equipment here, and requires such a person to nominate a local representative (Act 709, sections 2(2)(b), 2(3), 2(4)). The Consumer Protection Act 1999 takes the other route: where goods are manufactured offshore and the foreign manufacturer has no ordinary place of business here, the local importer or distributor is treated as the manufacturer and carries the liability (Act 599, section 3(1), definition of "manufacturer", paragraph (c)). We recommend one or both devices. A second limb emerged from the foreign-developer consultation at Annexure C, and we regard it as equally structural: where the Authority demands training data documentation, model documentation or assessment records from a foreign Developer, the Developer may be restrained by its home law (data-export approval regimes, state and business secrecy, confidentiality obligations) from producing them. A power vindicated only by penalising a party for obeying its own law is a standoff, not enforcement. The framework should state how such conflicts are handled, routing them through regulatory cooperation, a function the Authority already holds (Paper, p. 7), and which now has treaty-level content (¶19A), rather than leaving foreign Developers to choose which jurisdiction's order to breach. We note the Government has itself identified the mirror image of this problem: launching the Action Plan, the Prime Minister flagged the reach of foreign law such as the US CLOUD Act over data as a national concern (as reported, The Edge Malaysia, 28 July 2026). The collision of one jurisdiction's compelled disclosure with another's prohibitions is the same phenomenon in both directions, and a Bill that designs for it will be better placed to object to it.

F. Focus Area 3: AI Governance Principles

1. We support all five principles and the phased approach. Our comments go to how they connect to law that already binds Malaysian organisations; and, at paragraph 8 above, to the question Principle 5 currently presupposes.

2. **Principle 5 should integrate the PDPA, not gesture at it.** As drafted, data must be managed "in accordance with applicable laws". Everyone already had to do that. The drafting cost is that it treats Malaysia's most developed body of relevant law as background noise when it could be doing the Bill's work. The PDPA supplies today: seven statutory principles (section 5); data subject rights of access and correction; conditions

for processing sensitive personal data (section 40); a mandatory DPIA regime with defined triggers; and design-stage methodology. We recommend Principle 5, or its first implementing instrument, expressly provide that an AI system processing personal data discharges the data-governance principle by complying with the PDPA framework. Malaysian organisations then meet Principle 5 through programmes they already run, and the two regimes cannot drift apart because one incorporates the other.

3. **Principles 1 to 3 already have operative Malaysian content: the Bill should claim it.** Principle 1 provides that AI must not place important decisions entirely beyond human review; the ADMP Guideline states the operational version, that AI must not be relied upon as the sole factor in making decisions concerning a data subject (para 10.2.6). Principle 2 requires that affected persons be given suitable awareness; the ADMP Guideline supplies the mechanism, informing the data subject through the privacy notice that AI is used, in language neither excessively lengthy nor overly technical (para 10.2.3). Principle 3's "effective redress" maps onto the statutory rights of access and correction. We are careful not to overstate this: the ADMP Guideline frames these as best practices the data controller "may adopt" (para 10.2). That is the opportunity. The Bill can give statutory force to standards the Commissioner has already worked out and organisations have already seen, rather than drafting new formulations that will inevitably diverge. Cross-referencing costs a sentence; divergence costs a decade of parallel interpretation.

4. **On what Principle 5 presupposes.** For the reasons at paragraphs 8–10, the "properly sourced" and "provenance" language should be given content by guidance developed jointly with the Commissioner rather than defined unilaterally by the Authority, and the Bill should avoid drafting that assumes a lawful basis for training on personal data exists. We note this is a live constraint on the Bill's stated purpose of stimulating innovation (Paper, p. 3): an organisation that cannot identify a lawful basis will either not train in Malaysia, or will train without having identified one.

5. **A principle on fairness and non-discrimination should be added.** In answer to the Questionnaire's direct question: yes, one is missing, and the reason is domestic rather than imported. The Commissioner's own materials treat discriminatory algorithmic outcomes as a recognised harm: the ADMP Guideline's insurance-premium illustration warns that the approach "has the potential to lead to group discrimination among applicants", and its list of decisions with significant effects includes denial of employment opportunities, differential pricing, effects on access to education, and reputational harm (paras 7.4.3–7.4.6). Yet none of the five principles names fairness. This matters more in Malaysia than elsewhere because, as we develop at paragraph 34, discriminatory private-sector algorithmic outcomes will often contravene no written law here and so escape the risk framework as well. There is also a continuity point: the National Guidelines on AI Governance and Ethics (2024), the framework this Bill builds on, include fairness among their seven principles. Adopting it restores continuity rather than importing anything.

6. **"Due regard" needs a floor at the top tier.** The standard is right for low-risk systems. For high-risk systems a duty that is entirely self-assessed invites paper compliance: an organisation can honestly believe it had due regard and have nothing to show for it. For Tier 2 systems, due regard should be demonstrable: a documented assessment and an identified human-oversight checkpoint. The same discipline should apply to the oversight itself: a human "in the loop" who lacks the training, information or authority to question the system's output is presence, not oversight. We recommend the implementing instruments define oversight by capability: the person must understand what the system optimises for, have access to the information needed to evaluate its output, and be empowered to depart from it.

7. **Oversight should also be graduated, because the deployment pattern is changing.** The Bill reads as though a system produces one output informing one decision. Systems now arriving plan, call external tools, and take sequences of actions. The useful gradation runs from a human directing and approving every step, through approval at significant steps, to approval only at critical steps or on failure, to observation with audit

after the fact. That gradation has legal consequence in Malaysia, not merely operational consequence: in our view the last position cannot be reconciled with ADMP para 10.2.6 where the decision carries the significant effects the Guideline identifies, while the third can be, provided approval checkpoints sit where the action produces the legal or significant effect, and the approval is real rather than a formality. Two practical points follow for the Enablement function: regional material addressed specifically to this deployment pattern already exists and need not be written from scratch; and the implementing instruments should be able to recognise certification against an established AI management system standard, ISO/IEC 42001 being the obvious candidate, as evidence that the due-regard duty is discharged; the National AI Classification mechanism the Action Plan commits to at initiative E13 should be recognised on the same basis, so the Plan's certification and the Bill's compliance architecture reinforce rather than duplicate each other. One certification satisfying several regulators is a material reduction in compliance cost, and it costs the Authority nothing to accept evidence it did not have to generate.

G. Focus Area 4: AI Risk Framework

1. **First, a point in the framework's favour.** The fourth harm category, "contravention of any written law" (Paper, p. 11), already brings the PDPA inside the risk framework. An AI system likely to process personal data unlawfully is already a harm-relevant system on the Bill's own definition. We think this is correct and elegant, and we would encourage NAIO to say so expressly, because it changes the character of everything below: aligning the Bill's assessment machinery with the PDPA's is not an optional courtesy between regimes: the Bill has already made PDPA compliance part of its own harm test.

2. **The same category needs a gravity threshold.** Read literally, it makes every breach of every statute an AI harm: a technical licensing lapse registers on the same axis as a system that kills. A category with no floor does not make the framework stricter; it makes it unusable, because Developers and Deployers cannot triage against a category that everything satisfies. We recommend qualifying it: for instance, contraventions attracting criminal liability or significant regulatory penalty, or contraventions of laws protecting life, liberty, property or personal data. We flag this against our own interest in the point: having praised the category, we think the submission is more useful if it also says where the drafting will strain. The reasoning at paragraphs 8–9 makes the threshold more important, not less.

3. **The real gap: harm that contravenes no written law.** The four categories capture the catastrophic and the illegal. What escapes all four is the harm AI most commonly inflicts: outcomes that are lawful but damaging. Three examples, each from the Commissioner's own materials rather than foreign frameworks. *Financial harm*: the ADMP Guideline's first illustration is a fully automated credit-scoring system rejecting applications below a threshold "without human review", noting a single automated assessment may have lasting effects; the DPIA Guideline uses the same scenario as processing with "a direct and significant effect on the data subject's economic rights and opportunities". *Discriminatory treatment*: the insurance illustration identifies group discrimination as the risk, yet Malaysia has no general anti-discrimination statute reaching private-sector algorithmic decisions, so the very risk the Commissioner's guideline flags would not register as harm under the Bill. *Reputational harm and denial of access*: the ADMP Guideline's list of significant effects includes denial of employment opportunities, effects on access to education, and reputational harm. The Paper itself supplies the mechanism: the categories "may be expanded" (p. 11). We recommend expansion at the baseline rather than sector by sector, because these harms follow automated decision-making wherever it goes.

4. **The Tier 1 intent gate should be reconsidered.** Tier 1 captures systems "developed or deployed with an intent to cause harm". We understand the instinct, but as a gate it points the wrong way. A system

catastrophically dangerous by design or by indifference escapes Tier 1 wherever intent cannot be proved, which in practice is almost always. Meanwhile a system genuinely deployed with intent to cause the defined harms is already criminal conduct, and freshly so: the Cybercrimes Bill 2026 criminalises precisely the intent cases, including transmitting deceptive synthetic content with intent to facilitate an offence (clause 23), computer-related fraud (clause 17) and identity theft by means of a computer system (clause 22). The result is a top tier simultaneously too narrow to catch what it should and largely redundant for what it does. We recommend Tier 1 be defined by **use-case severity**, categories whose risk is unacceptable regardless of the operator's state of mind, with intent operating as an aggravating factor in enforcement. Since the Questionnaire invites benchmarks, one is directly on point: the EU AI Act prohibits these practices by their objective *or their effect* (Article 5(1)), and its recitals state expressly that it is not necessary for the provider or deployer to have intended significant harm (Recital 29).

5. **The three classification axes will conflict, and the Bill should say which governs.** The Paper proposes requirements be issued by class of AI system (Developer as primary duty holder), by sector, and by domain of deployment (Deployer as primary duty holder) (p. 12). Each is sensible alone; together, without a primacy rule, they produce two problems in ordinary configurations. **The tier becomes indeterminate**: a general-purpose model may be unremarkable by class and unacceptable by domain, the same model behind a marketing tool and behind clinical triage. **And the duty holder becomes uncertain**: where one axis names the Developer and another the Deployer, each can reasonably read the framework as placing the obligation on the other, and the gap is discovered only after something goes wrong. We recommend the Bill state that where more than one axis applies, **the highest applicable classification governs**, and that Developer-side and Deployer-side duties are cumulative rather than alternative, with the Malaysian-established Deployer as compliance interface where the Developer is offshore. This is consistent with the Paper's foundation of allocating accountability by degree of control; it makes explicit that control is layered rather than exclusive. A related observation: if Tier 1 is confined to proven intent and Tier 3 is residual, almost every system of regulatory interest lands in Tier 2 and the framework is functionally binary, which argues for graduated obligations *within* Tier 2, scaled to the Paper's own evaluation factors.

6. **One assessment, not two.** Tier 2 systems will face "risk assessment, documentation, internal controls, human oversight, monitoring, and mitigation measures". For any AI system processing personal data at scale, this describes an obligation Malaysian organisations already carry. The DPIA regime is mandatory where quantitative thresholds are met (more than 20,000 data subjects, or more than 10,000 where sensitive personal data including financial information is involved) and below those, where the Data Protection Officer's assessment of qualitative factors requires it. Those factors read like a checklist written for the Bill: automated decision-making and profiling posing high risk; innovative technology; systematic monitoring; targeting of children or vulnerable individuals; potential legal or significant effects. The DPIA methodology scores likelihood against impact; the Bill proposes likelihood, severity and scale, and duration and reversibility. These are the same exercise under two names.

7. Running it twice (once for the Commissioner, once for the Authority, on the same system, to answer the same question) is deadweight cost felt most by the SMEs the Paper says it wants to protect. We recommend a **mutual-recognition rule**: where an AI system processes personal data, a DPIA conducted in accordance with the Commissioner's guideline satisfies the Bill's Tier 2 risk-assessment obligation in respect of the personal-data dimensions, and an assessment under the Bill satisfies the DPIA requirement to the corresponding extent. The DPIA Guideline has effectively left the door open from its side: it does not derogate from requirements in other instruments, and where a broader obligation exists the broader requirement applies (para 7.3). The Bill should walk through it from this side. A final architectural point: the Commissioner's framework already names the internal judgment-holder: the mandatorily appointed Data Protection Officer decides whether qualitative factors trigger a DPIA. For AI systems processing personal data, that officer is

the natural owner of the Bill's risk classification, and recognising this avoids forcing organisations to invent a parallel "AI officer" role duplicating a function they were required to staff last year.

H. Focus Area 5: AI Incident Reporting

1. We support a national framework, and specifically support the inclusion of near misses, early-warning value is where incident regimes earn their cost. Our recommendations address an overlap that is not two-way but three-way.

2. **A large class of AI incidents are also personal data breaches**, triggering the mandatory regime operationalised in Circular No. 1/2025: notification to the Commissioner with prescribed particulars within 72 hours; notification to affected data subjects within 7 days where the breach causes or is likely to cause significant harm; staged provision of information completed within 30 days; records retained at least two years. **For critical sectors, the same event triggers a third regime.** Under section 23(1) of Act 854, a national critical information infrastructure entity that knows a cyber security incident "has or might have occurred" must notify both the Chief Executive of NACSA and its sector lead within the prescribed period; non-compliance carries a fine up to five hundred thousand ringgit or imprisonment up to ten years. The designated entities (banking and finance, healthcare, energy, transport, telecommunications and government among the scheduled sectors) are precisely the organisations deploying the most consequential AI. A serious AI incident at a Malaysian bank could engage three reporting duties on three clocks to three regulators. A fourth dimension may overlay them: where the incident involves criminal conduct, the Cybercrimes Bill 2026 brings investigation, search and seizure powers and its own coordination body, so the Authority's technical fact-finding may run alongside a criminal investigation into the same systems, making evidence preservation and primacy rules part of what the coordination protocol must cover. We add a terminological observation: Act 854 also builds its architecture on "sector leads" (sections 15–16), so the coordination problem already has a statutory sibling whose vocabulary the Bill partly shares.

3. The Paper states the framework "leverages existing reporting mechanisms". We recommend making that concrete in the two places it is needed on day one: **designate the Commissioner's breach-notification regime and the Act 854 regime as recognised existing mechanisms, and provide for single-window filing**: one report, in a format satisfying each engaged regime, shared between the Authority, the Commissioner and NACSA. Without this, an organisation in the first 72 hours of a serious incident will be preparing two or three overlapping reports on different clocks, and the incident response itself will suffer. Timelines should be specified, and specified compatibly with the 72-hour / 7-day / 30-day structure and the prescribed periods under Act 854. **The trigger should be knowledge, not occurrence.** A duty to report assumes the ability to detect, and detection of harmful AI behaviour is genuinely hard. An occurrence-triggered duty would impose strict liability for failing to report incidents the organisation never saw, punishing incapability rather than concealment and giving sophisticated actors an incentive not to look. Act 854 draws the line correctly, and the Bill should adopt the same knowledge-based trigger, paired with the Enablement function building detection capability. **The duty should also be layered, and the first filer named.** Developed jointly with the foreign-developer consultation at Annexure C, we recommend a three-layer structure that is cleaner than the Paper's single category: *incidents*, events causing or reasonably likely to cause the defined harms, reportable mandatorily on the knowledge trigger; *near misses* reportable voluntarily with the safe-harbour protection described at paragraph 42, preserving the early-warning value the Paper rightly wants; and *mere weaknesses*, no harm occurred or reasonably foreseeable, attracting a duty to mitigate, not to file, since complex systems inevitably contain weaknesses that never mature into harm. And because determining which party bears underlying responsibility can require root-cause investigation and contract review, the framework should designate the **Deployer as primary reporting party** (the party with

operational visibility, and typically the Malaysian-established one) with allocation between Developer and Deployer resolved after filing, not before it. A reporting duty that waits for the parties to agree who is responsible will always file late.

4. **What the report should contain.** Single-window filing only works if the Bill's content set is a **superset** of what the other regimes require; otherwise "one report" means one report plus a second for the fields the first omitted. Circular No. 1/2025 already prescribes, for the data-breach dimension: the date and time the breach came to the controller's knowledge; the type of personal data and type of breach; the method of identification and believed cause; the number of affected data subjects and estimated records; the systems affected; likely consequences; the chronology leading to loss of control; remedial and mitigating steps; steps in respect of affected data subjects; and a contact point. We recommend adopting that structure and adding the four elements specific to AI: **the system and version that acted**, uniquely identified, without which no chain of accountability survives a model update; **the risk classification and whether an assessment existed**, with the DPIA reference where applicable; **whether human oversight was designed in and whether it was exercised**; and **whether an automated decision affecting an identified person, or the generation of synthetic content, was involved.** Reusing the Circular's field structure has a second benefit: organisations and their advisers have been building processes against those fields since 2025, and a form that reuses them will be completed accurately in the first 72 hours. A newly drafted form asking the same questions in different words will not. Two further design points: near-miss reporting only works if a good-faith report does not become evidence against the reporter, so a limited safe harbour is needed; and the public-complaint channel should connect to the data subject rights regime, so one complaint by an affected data subject can engage both processes rather than requiring an ordinary person to work out which regime their problem belongs to.

I. Focus Area 6: AI Sandbox

1. We support the sandbox, including implementation through Sectoral Leads' existing infrastructure. **The PDPA applies inside it, and the framework should say how.** The Paper's own list of functions includes "facilitating controlled access to data by anonymising or synthesising personal information" (p. 16), an acknowledgment that real personal data will move through sandbox environments. Nothing in Act 709 switches off because testing is supervised. Rather than leaving each cohort to work this out, the framework should build it in: a DPIA at entry for cohorts processing personal data in ways likely to meet the triggers; data minimisation and purpose limitation as standing testing conditions; breach-notification obligations expressly unaffected; and anonymisation and synthetic-data standards set with the Commissioner's input. Framed this way the requirement is a benefit: a participant exits with a validated data-protection posture alongside its AI assessment, which is worth something commercially.

2. **Exit should carry defined regulatory effect.** The Questionnaire itself identifies "regulatory uncertainty post-sandbox" as an anticipated barrier. The fix belongs at design stage: successful exit should produce something with legal weight (a provisional risk classification, or a time-limited safe harbour for the system as tested) so participation reduces the participant's regulatory risk rather than merely informing the regulator's policy. If exit produces only a report, the sandbox will attract those who least need it and be avoided by the startups it is designed for.

44A. **The sandbox should also be framed as what it can become: a regional entry pathway.** The Paper presents the sandbox around Malaysian developers, start-ups and SMEs (p. 15). The foreign-developer consultation at Annexure C confirms a second market for it: developers entering ASEAN would use a Malaysian sandbox as a supervised soft-landing, *if the exit outcome carries weight, and travels*. Both conditions have proven designs. On

weight, the EU AI Act requires exit reports to be taken positively into account in subsequent conformity assessment, accelerating it to a reasonable extent (Article 57(7)), the legal-effect-on-exit principle at ¶44. On travel, the EU already operates mutual recognition *internally*: participation in a sandbox established by one Member State is mutually and uniformly recognised and carries the same legal effects across the Union (Article 58(2)(g)). We recommend Malaysia propose the ASEAN application of that design, mutual recognition of sandbox outcomes across ASEAN jurisdictions, and no member state is better placed to table it: Malaysia holds the ASEAN AI Safety Network secretariat, is a founding member of the World AI Cooperation Organization (¶19A), and would operate one of the region's first statutory AI sandboxes. The venue for tabling it already exists in the Government's own plans: the Action Plan commits Malaysia to lead the ASEAN AI Safety Network in aligning regional safety and security standards and to convene Malaysia-chaired regional networks of regulators (initiative E12). Mutual recognition of sandbox outcomes is the natural first deliverable of exactly that alignment work. A sandbox whose exit certificate is respected across ASEAN is a stronger inbound-investment draw than any incentive currently in the Paper; the SME measures the EU pairs with its sandboxes (free access (Article 58(2)(d)), priority access (Article 62(1)), outreach and pre-deployment guidance) complete the package for the domestic cohort.

1. **A drafting observation.** We read the Paper's list of functions as describing two different instruments. Eight of the nine (isolated execution, non-production development, integration testing, controlled data access, model evaluation, limited-user testing, workflow trialling, simulated environments) are functions of a *technology* sandbox: a technical facility requiring no statutory power beyond the authority to operate it, altering nobody's legal obligations. The ninth is different in kind: permitting "regulated activities or innovations to be tested under a relaxed regulatory environment" (function (i), p. 16) is a *regulatory* sandbox, requiring an express statutory power to modify or disapply requirements, a defined scope of what may be relaxed, and a rule about what cannot be. The distinction may matter for drafting because of that third element: **the Bill cannot relax obligations arising under other statutes.** A participant remains bound by Act 709 in full unless the Bill expressly provides otherwise and the PDPA framework accommodates it. Unless the two are separated, participants may assume "relaxed regulatory environment" means relief from data-protection obligations, and the Commissioner's enforcement position will be unaffected by their assumption.

J. Cross-border data flows

1. AI development and deployment routinely move Malaysian personal data offshore: for training, for inference, for hosting. There may be pressure on the Bill to answer this with a new data-localisation or sovereignty standard. We suggest it should not, because Malaysia already has a considered answer: the cross-border transfer regime under Act 709 and the Commissioner's Cross-Border Personal Data Transfer Guideline, under which a data controller conducts a structured Transfer Impact Assessment of the receiving jurisdiction (examining whether it provides comparable data subject rights, security obligations, breach-notification requirements and an enforcing regulator (para 5.4)) with findings valid for no more than three years (para 5.6). A single provision confirming that offshore AI data flows are governed by that regime would prevent two competing Malaysian transfer standards, which would serve no one.

---

1. Part One deals with the Bill's internal design. **Annexure B deals with its external fit**, and we summarise its findings here because they do not map onto any of the Paper's guiding questions and would otherwise be easy to miss. Four statutes the consultation documents do not discuss will bear directly on how the Bill works, and in each case our point is not that the Bill should absorb these regimes, most of the reforms they may eventually need belong in their own statutes, but that the Bill should be drafted with knowledge of them.

2. In summary:

3. **The Online Safety Act 2025.** Generative AI is now the dominant production technology for several categories of harmful content the Act addresses, and deepfake investment scams are the most visible AI harm affecting ordinary Malaysians today. Act 866 reaches the platform that *carries* 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. We recommend three connections, including consideration of MCMC as Sectoral Lead for AI-generated content harms, and a routing rule between the two public-complaint channels. **We also recommend a proportionate synthetic-content transparency obligation**: the missing joint on which the platform detection duties, the criminal provisions and the Bill's own incident channel all depend, since no rule currently requires an AI system's output to be identifiable as synthetic.

4. **The Consumer Protection Act 1999.** Principle 3 promises "effective redress", but the Paper is silent on whether a person injured by an AI system can recover from anyone. Regulatory sanction and civil compensation are different things. We recommend the Bill state its position, whether breach of a duty under it is actionable, because silence will be litigated and will be priced into every AI supply contract in Malaysia from the day it takes effect.

5. **The Evidence Act 1950.** Section 90A presumes a computer producing documents in the course of its ordinary use is operating properly. Generative AI breaks that assumption's logic without breaking its language: a model producing a confident falsehood is not malfunctioning. We recommend the Bill's documentation and traceability duties be designed with their evidential destination in mind.

6. **The Copyright Act 1987.** Act 332 contains no exception addressed to text and data mining, and no provision for works generated without a human author. Principle 5's "properly sourced" language presupposes answers on the input side that copyright law has not given, the same structural problem as Finding 2, in a different register. We recommend coordination and restraint.

K. Closing

1. The theme is deliberately consistent across all four documents. Malaysia finished building a data-governance framework for AI in 2025; it has an operating cyber incident regime, a new online content regime, a product liability statute, an evidence code and a copyright act, each of which the Bill will touch on its first day in force. What the Bill needs to do, and what only the Bill can do, is connect: one vocabulary for the actors, one assessment per system, one filing per incident, one transfer standard, a fairness principle to catch the harms that fall between the statutes, a stated position on civil redress, and drafting that neither collides with the evidence and copyright frameworks nor quietly presupposes their reform. Each is a choice available now, at the pre-drafting stage, and each becomes progressively harder to retrofit once the regimes develop separate practice.

2. We would be glad to discuss any part of this submission with the AI Policy Team, and to assist further as the Bill develops.

Yours faithfully,

Edwin Lee Founder and Managing Partner **Edwin Lee & Partners** edwin@lpplaw.my | +6011-5954 1201 | lpplaw.my 31 July 2026

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.