ANNEXURE C

Annexure C: The Bill as Experienced by a Foreign Developer Entering Malaysia

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

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

C1. Why this Annexure exists

C1.1 Malaysia's stated ambition is to attract AI investment and to become a regional AI hub, an ambition underscored, since the Paper was issued, by Malaysia's founding membership of the World AI Cooperation Organization on 16 July 2026 (core submission, ¶19A). A significant proportion of the AI systems that will be placed on the Malaysian market, or used by Malaysian Deployers, will be developed outside Malaysia.

C1.2 Those developers are the parties least likely to respond to this consultation. The reasons are practical rather than indifferent: language, unfamiliarity with the Malaysian consultation process, short timelines, and limited incentive to engage with a foreign framework at pre-drafting stage. The consequence is that the Bill risks being drafted without direct evidence of how its central definitions (Developer, Deployer, the extraterritorial trigger, the risk tiers) will land on the organisations it most needs to reach.

C1.3 This Annexure puts some of that evidence before NAIO. We posed five questions to practitioners who advise AI developers active in this region, and received substantive written responses, which we present below in our structure with the contributor's views attributed throughout and our own observations following each section. We offer this material on the footing that the perspective is otherwise unlikely to be before the AI Policy Team.

C2. Method, scope and attribution

C2.1 The contributor. The responses presented in this Annexure were contributed by **Global Law Office ("GLO")**, established in 1984 as the first law firm founded following the implementation of China's reform and opening-up policy and today among the largest full-service law firms in the People's Republic of China, whose practice includes advising Chinese and international AI and technology companies. The responses were prepared by **Xu Guosheng** and **Dai Chang** of GLO, whose contribution is gratefully acknowledged; the firm is introduced more fully at paragraph 2A of the core submission. References in this Annexure to "the contributor" are references to GLO.

C2.2 Scope of what is being said, and by whom. The contributor is not advising on Malaysian law, and Edwin Lee & Partners is not advising on the law of the People's Republic of China or any other jurisdiction. Where this Annexure describes foreign legal requirements or foreign market practice, those descriptions are the contributor's and are provided as background to the Malaysian questions; the contributor's citations are preserved and marked as such in the notes at the end of this Annexure. Where this Annexure states conclusions about the operation of the proposed Bill, or refers to the EU Artificial Intelligence Act, those statements are ours: we have independently verified each reference to the EU AI Act in this Annexure against the text of Regulation (EU) 2024/1689.

C2.3 A note on framing. The perspective set out here is that of a *foreign developer* generally. The illustrations are drawn principally from the Chinese market because that is where the relevant deployment activity and the contributor's experience are concentrated, not because the analysis is confined to developers from any one country. The same questions arise, in substantially the same form, for a developer established in any jurisdiction outside Malaysia.

C3. Scope and definitions: who is the Developer?

C3.1 The question we posed. The Bill defines a Developer as any person or organisation that materially shapes what an AI System is capable of doing, expressly including a party that adapts a model for a specific use case, integrates it into a wider system, or later modifies it after deployment (Paper, p. 5). We asked: where a foreign developer releases a model under open weights and a Malaysian company downloads and fine-tunes it, does the Malaysian company become a Developer alongside the original creator, and is that dual-Developer outcome workable? Where a model is supplied purely via API, at what point does the customer's adaptation cross into "materially shaping"? And how does the definition compare with the regimes a foreign developer already operates under?

C3.2 The contributor's response. The contributor observes that multiple parties assuming Developer-type responsibilities is not uncommon, particularly where foundation models, fine-tuned models and layered AI applications are involved, and notes that the EU AI Act addresses this expressly: under Article 25, a distributor, importer, deployer or other third party is treated as the provider of a high-risk AI system in defined circumstances, including where it makes a substantial modification to such a system.[1]

The contributor draws a contrast with the position at home: China's Interim Measures for the Management of Generative Artificial Intelligence Services impose obligations primarily on entities *providing generative AI services to the public within China*. Where an entity does not provide a public-facing service, the mere release of open-source model weights, model architecture or training information would generally not, of itself, make that entity a service provider under the Measures.[2] The contributor observes that, under the definition as proposed, providers of open-source models may thereby assume Developer obligations under the Bill.

The contributor makes two recommendations to the Malaysian drafters. First, establish a clear responsibility-allocation framework for the open-source ecosystem: foundation model providers, model modifiers and application developers exert different degrees of influence over an AI System, and allocation should follow each actor's business model, technical capability and actual degree of influence at the relevant lifecycle stage. Second, clarify the criteria for "materially shape", and in particular whether the following factors are to be taken into account: whether the actor modifies model weights; whether it conducts additional training such as fine-tuning; whether it changes the model's underlying capabilities; whether it affects the system's output behaviour; and whether it affects the system's safety mechanisms or access-permission settings.

C3.3 Our observations. Three points.

- **The five-factor list is drafting-ready, and we commend it to NAIO with one refinement.** The factors fall into two natural bands: the first three (weights, further training, capability change) alter what the model *is*, while the last two (output behaviour, safety and permission settings) can be effected by configuration alone: a system prompt affects output behaviour. We suggest the implementing instruments treat the first band as presumptively "materially shaping" and the second as shaping only where the effect is substantial and persistent. That preserves the Bill's intent while keeping ordinary enterprise configuration on the Deployer side of the line: the answer to our own question C3.1(b).
- **Both major regimes accommodate open-source release; the Bill as drafted accommodates neither's approach.** The contributor's account of the Chinese position has a directly comparable EU analogue we verified independently: the EU AI Act does not apply to AI systems released under free and open-source licences unless they are placed on the market as high-risk or fall within specified categories (Article 2(12)), and its recitals encourage rather than mandate documentation practices for open-source components (Recital 89). The Bill's Developer definition, by contrast, attaches at the moment of shaping, with no accommodation for the release model. We do not suggest an open-source exemption, an open-weights model can do everything a hosted one can, but the dual-Developer outcome needs the allocation rule the contributor asks for, and the

core submission supplies its Malaysian mechanics: cumulative duties allocated by the layer each party controls (R15), an express controller/processor mapping (R3), and a substantial-modification threshold (R8). Article 25(2) of the EU AI Act adds a further element worth borrowing: when a downstream party becomes the provider through modification, the initial provider must cooperate and hand over the information and technical access the new provider needs. Without a hand-over duty of that kind, the Malaysian fine-tuner becomes a Developer on paper while lacking the documentation to discharge the role in fact.

- **The structural consequence, in our reading, is an enablement point.** Putting the contributor's two propositions together: a provider whose home regime attaches obligations to the provision of a public-facing service, and attaches none to the release of weights as such, meets in this Bill a Developer definition that attaches to the act of shaping the system itself. The obligation is unfamiliar in structure, not merely in detail. That inference is ours, not the contributor's, but it follows directly, and its consequence is practical: if the Bill's role structure differs from what the largest inbound developers know from home, the Authority's guidance for foreign Developers is not a courtesy; it is a condition of the extraterritorial scope functioning at all.

C4. Extraterritorial reach

C4.1 The question we posed. The Bill covers AI systems used by a Deployer established in Malaysia regardless of where the system is hosted (Paper, p. 5). We asked how that trigger compares with what a foreign developer already manages elsewhere, whether it is workable, and, pointedly, whether the enforcement anchor we recommend at core ¶24 (a local representative for foreign Developers of high-risk systems, and/or the Malaysian Deployer as express compliance interface) would deter foreign developers from serving this market.

C4.2 The contributor's response. The contributor regards the proposed scope as broadly aligned with emerging international approaches, particularly the EU AI Act's extraterritorial coverage of systems placed on the market, put into service or used within the jurisdiction, and notes that China's generative AI framework is service-oriented, regulating services provided to users within China, without an equivalent comprehensive extraterritorial framework for foreign providers.[3]

The contributor's view is that an extensive territorial scope should be supported by three things: a clear and comprehensive set of compliance obligations for the Developers and Deployers it reaches; phased implementation with sufficient transition periods for domestic and foreign stakeholders to build compliance mechanisms; and, the point the contributor develops furthest, **regulatory coordination with other jurisdictions on cross-border information disclosure**. Where a foreign AI provider is required to produce training data documentation, model documentation, training details or assessment records, the contributor cautions that Malaysian authorities should consider potential conflicts with the provider's foreign data protection obligations, business-secret protections, confidentiality obligations and cross-border data transfer requirements.

C4.3 Our observations. The disclosure-conflict point is, in our view, the single most valuable contribution in this Annexure, and we have promoted it into the core submission (¶24; R7 as extended). The problem it identifies is structural: the Authority's investigation and information-gathering powers assume the addressee is *able* to comply, but a foreign Developer may be restrained by its home law (data-export approval regimes, secrecy protections, confidentiality obligations) from producing the very documentation the Authority demands. A power that can only be vindicated by penalising a party for obeying its own law is not an enforcement mechanism; it is a standoff. The answer is not to weaken the power but to route the hard cases through regulatory cooperation, and the Bill's own architecture already contains the vehicle, since international technical cooperation is among the Authority's stated functions (Paper, p. 7) and now has treaty-level content (core ¶19A).

We record one further matter candidly. The contributor did not comment on the local-representative and Deployer-as-interface devices we specifically asked about. We do not treat that silence as endorsement; we note only that a contributor advising the companies most affected raised no objection to devices that the closest comparator regime, Korea's AI Basic Act, already imposes, and that the contributor's own asks (clear obligations, transition periods, coordination) are conditions of workability rather than arguments against reach. Recommendation R7 therefore stands as made, now carrying both limbs: the anchor, and the disclosure-conflict mechanism.

C5. Risk tiers in practice

C5.1 The question we posed. Which common categories of foreign AI products entering this market would fall into the high-risk tier as described: recommendation engines, generative assistants, credit and fraud scoring, smart vehicle systems, consumer-device AI? Would that classification surprise the companies behind them, and is the harm-anchored framing clear enough for a foreign developer to self-classify without local advice?

C5.2 The contributor's response. The contributor anticipates that e-commerce credit and fraud scoring, smart vehicle systems, and AI features in consumer devices will be placed in Tier 2, and states that confirmation of this classification, and of the treatment of other common AI applications, would be appreciated. The contributor is concerned that the inclusion of "contravention of any written law" as a baseline harm category "brings significant uncertainty, and unnecessarily broadens the scope of the framework and scope of supervision." The contributor adds that operators active across regions will value consistent definitions and risk tiering, since users of their systems may otherwise face burdens of demonstrating equivalence between jurisdictions' standards.

C5.3 Our observations. Three points, the second of which we think NAIO will find useful.

- **The request for confirmation is itself the answer to our self-classification question.** A sophisticated firm advising the largest AI companies in the region read the framework and asked the regulator to confirm the tiering. If self-classification were achievable from the Paper's four harm categories, the request would have been unnecessary. This is direct evidence for the defined high-risk threshold at R8 and for our test of workability: a threshold a compliance officer can apply without external advice (core ¶23).
- **The contributor's two observations, taken together, make the case for R13 more completely than either does alone.** The contributor expects credit and fraud scoring to be Tier 2, and separately asks that the "contravention of any written law" category be narrowed. Under the four categories as drafted, those two outcomes are difficult to achieve simultaneously: credit scoring reaches Tier 2 *only through* the fourth category (via the PDPA and ADMP dimensions of the processing) (it causes no death, no bodily injury, no deprivation of liberty), so narrowing that category without more would move credit scoring to Tier 3. What delivers both of the contributor's outcomes at once is precisely R13: name financial harm, unjustified discrimination and denial of access as harm categories in their own right, so that credit scoring is Tier 2 on its merits, and then apply a gravity threshold to the fourth category so it retains its integrating function without sweeping in every technical contravention. Two perspectives, one drafting answer.
- **The categories the contributor did not address are the ones that expose the gap most sharply.** We asked about recommendation engines and generative assistants; the response addresses neither. Under the four categories as drafted, both classify as Tier 3: outcomes we analyse at core ¶34 as the framework's principal blind spot, since the harms those systems most plausibly cause (financial, discriminatory, reputational) contravene no written law and appear in no category.

C6. Incident reporting

C6.1 The question we posed. Would a separate Malaysian AI incident reporting channel be manageable for a foreign developer already carrying incident and filing obligations at home, and what would make it workable, including whether the single-window arrangement we recommend at R16 would make a material difference?

C6.2 The contributor's response. The contributor seeks clarification on how the proposed mechanism interacts with existing sectoral mechanisms: whether "leveraging" existing mechanisms means reporting is conducted *through* the relevant Sectoral Lead or in *parallel* to both the Sectoral Lead and the Central AI Authority; how inconsistencies between different regulators' findings or recommended mitigations would be resolved; whether reporting must occur within a prescribed timeline or on a best-effort basis; and whether an existing sectoral deadline also governs under the AI framework.

On the definition of the reportable event, the contributor suggests distinguishing actual AI incidents from general system weaknesses: where an organisation identifies a weakness but no harm has occurred or is reasonably foreseeable, the primary obligation should be to take commercially reasonable steps to mitigate the risk, rather than mandatory reporting, since complex systems inevitably contain weaknesses that never result in harm.

On who reports, the contributor observes that determining which party has control and responsibility may require root-cause investigation and review of the contractual arrangements between Developer and Deployer; requiring the parties to agree responsibility before filing would delay reporting. For comparison, the contributor notes that in China a report is required within four hours for "relatively major" cybersecurity incidents, a full report within 72 hours, and a post-incident analysis within 30 days.[4] The contributor recommends designating a clear primary reporting party, with responsibility allocation between the parties addressed subsequently, and suggests that consideration be given to placing the primary reporting obligation on the Deployer, as the party operating the AI system and having direct visibility of its practical impact. The contributor also suggests the framework clarify the position where harm arises not from the AI system but from supporting infrastructure such as networks, in which case the reporting obligation should sit with the network operator; and proposes that the Authority establish accessible public channels (hotlines, online enquiry channels) for complaints and enquiries.

C6.3 Our observations. The contributor's coordination questions are, almost verbatim, the questions our core submission answers at ¶40–41, asked here by the parties who would have to live under the regime, which we offer to NAIO as evidence that the single-window design at R16 addresses a real rather than theoretical concern. Two of the contributor's points improve our recommendation, and we have adopted both:

- **The three-layer structure.** Combining the contributor's incident/weakness distinction with the near-miss safe harbour we had proposed produces a cleaner design than either alone, and cleaner than the Paper's: **incidents** are reportable mandatorily on a knowledge-based trigger; **near misses** are reportable voluntarily with safe-harbour protection, preserving the early-warning value the Paper rightly wants; and **mere weaknesses**, no harm occurred or reasonably foreseeable, attract a mitigation duty, not a filing. R16 now carries this structure.
- **Deployer-first filing.** A reporting duty that waits for Developer and Deployer to agree who is responsible will always file late. Designating the Deployer as primary reporting party (the party with operational visibility, and typically the Malaysian-established one (core ¶24, R7)) with allocation of underlying responsibility resolved afterwards, gets the report in within the timelines that matter. This also coheres with the cumulative-duties rule at R15.

The comparative datum is worth a sentence of its own: operators in the contributor's home market already report within **four hours** for the most serious category. A Malaysian single window aligned to the Circular No. 1/2025 structure (72 hours to the regulator, staged completion within 30 days (core ¶40–42)) is not a burden by the standards these companies already meet; it is hospitable. The infrastructure point maps in Malaysia onto the network-operator

regimes under the Communications and Multimedia Act 1998 and, for critical sectors, Act 854, a boundary the coordination protocol at R1 should state rather than leave to inference.

C7. Sandbox access

C7.1 The question we posed. Would a foreign developer realistically use a Malaysian AI sandbox as a soft-landing route into the region, and what conditions would make participation attractive?

C7.2 The contributor's response. Four proposals. First, **cross-agency coordination**: AI risks extend beyond technical safety into competition, unfair commercial practices and consumer protection, so the sandbox framework should enable relevant authorities beyond the Sectoral Lead to guide or participate where their expertise is engaged. Second, **SME support**, for which the contributor cites the EU AI Act's measures (free participation for SMEs and start-ups, priority access, dedicated outreach, pre-deployment guidance) and recommends the Malaysian equivalents. Third, **recognition of sandbox outcomes**: the contributor notes that under the EU AI Act, exit reports are to be positively taken into account in subsequent conformity assessment and may accelerate it,[5] and that in Beijing, enterprises completing sandbox testing and obtaining an exit certificate may receive dedicated innovation support,[6] and recommends that Malaysian exit reports be positively considered in subsequent classification, assessment, certification or authorisation, adding that, given Malaysia's role in the regional AI ecosystem, consideration should be given to **mechanisms for broader recognition of sandbox outcomes across ASEAN jurisdictions**, which would enhance the sandbox's attractiveness for companies seeking regional expansion. Fourth, **safe harbour**: a general reference to a "relaxed regulatory environment" does not give a participant certainty about its position if unforeseen issues arise; where a participant acts in good faith within a sandbox plan approved by the Authority, administrative penalties should not be imposed solely in relation to activities within the approved scope.[7]

C7.3 Our observations. This is the affirmative answer to our question: a foreign developer would use the Malaysian sandbox as a regional entry route *if the exit outcome carries weight, and travels*. Three of the four proposals converge with positions already in the core submission: the safe harbour and exit-effect points are ¶44–45 and R17, and we have verified the contributor's EU citations against the Regulation (the exit-report provision is Article 57(7); the SME measures are Articles 62(1) and 58(2)(d); the fine-protection provision is Article 57(12)). The cross-agency point extends our Commissioner co-supervision recommendation naturally, and connects to the consumer-protection analysis at Annexure B Part 2.

The fourth proposal is new, and we have adopted it as **Recommendation R24**, with one addition the contributor's material did not include and which we located independently: the EU has already legislated exactly this mechanism *internally*: under Article 58(2)(g) of the EU AI Act, participation in a sandbox established by one Member State "is mutually and uniformly recognised and carries the same legal effects across the Union." ASEAN mutual recognition of sandbox outcomes is therefore not an experiment; it is the regional application of a design already operating in the world's most developed AI regime. And Malaysia is singularly placed to propose it: it holds the ASEAN AI Safety Network secretariat, it is a founding member of the World AI Cooperation Organization, and, if this Bill proceeds, it will operate one of the region's first statutory AI sandboxes. The Paper frames the sandbox around Malaysian developers, start-ups and SMEs (p. 15); we suggest it be framed, additionally and explicitly, as an **inbound-investment instrument and regional entry pathway**. For a country pursuing an AI investment strategy, a sandbox whose exit certificate is respected across ASEAN is a stronger draw than any incentive in the Paper as it stands.

C8. Summary of findings

C8.1 Five findings from this exercise, in descending order of consequence:

1. **The Authority's information demands on foreign Developers will collide with foreign law**, and the framework needs a stated mechanism (regulatory cooperation, not penalty for legally-compelled non-production) for that collision (adopted into core ¶24; R7 as extended).
2. **A Malaysian sandbox becomes a regional entry pathway if its exit outcomes are recognised across ASEAN**, a mechanism the EU already operates internally under Article 58(2)(g), and Malaysia is uniquely positioned to propose it (new R24).
3. **The reporting duty should be three-layered and Deployer-first**: mandatory for incidents, voluntary and protected for near misses, a mitigation duty for weaknesses, with the Deployer filing first and responsibility allocated afterwards (adopted into R16).
4. **Foreign operators cannot self-classify under the four harm categories as drafted**, evidenced by the contributor's own request for confirmation of the tiering, and the contributor's positions on Tier 2 and on the fourth harm category are reconcilable only through the expanded taxonomy and gravity threshold at R13.
5. **The Developer definition is unfamiliar in structure to the open-source providers it most needs to reach**, in whose home regime a weight release attracts no service-provider obligations; the five-factor "materially shape" list offered by the contributor, banded as we suggest at C3.3, is ready material for the implementing instruments (supporting R3, R8 and R15).

C8.2 These findings are reflected in the core submission at ¶12, ¶24 and ¶44A, and in recommendations R7, R13, R15, R16 and R24.

Contributor's citations

The following notes carry the contributor's citations. The references to instruments of the People's Republic of China are the contributor's and are reproduced as provided. Edwin Lee & Partners has independently verified the European Union references against Regulation (EU) 2024/1689, and where verification indicated a more precise rendering of the Regulation's text, the note below reflects the verified wording.

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. Statements as to the law or market practice of any jurisdiction other than Malaysia are attributed to the contributor identified at C2.1 and are provided as background only. Specific advice should be sought for particular circumstances.

---

1. Regulation (EU) 2024/1689 (Artificial Intelligence Act), Art. 25(1): distributors, importers, deployers or other third parties may be considered providers of high-risk AI systems where, among other circumstances, they substantially modify a high-risk AI system or place it on the market under their own name or trademark. ↩
2. Interim Measures for the Management of Generative Artificial Intelligence Services (生成式人工智能服务管理暂行办法), Art. 2: the Measures apply to organisations and individuals providing generative AI services to the public within the territory of the People's Republic of China. ↩
3. As above, Art. 2 of the Interim Measures. ↩

4. Measures for the Administration of Cybersecurity Incident Reporting (国家网络安全事件报告管理办法), Arts. 4 and 8. ↩

5. Regulation (EU) 2024/1689, Art. 57(7): the competent authority shall provide an exit report, and exit reports and written proof shall be taken positively into account by market surveillance authorities and notified bodies with a view to accelerating conformity assessment procedures to a reasonable extent. ↩

6. Beijing Government, 关于发挥北京数据基础制度先行区优势，加快打造数据基础制度综合改革试验区的实施方案 (2025): enterprises completing data regulatory sandbox testing and obtaining an exit certificate may receive dedicated innovation support services. Refer to: https://www.beijing.gov.cn/zhengce/zhengcefagui/202510/t20251029_4243376.html ↩

7. Beijing Economic-Technological Development Area, 关于加快推进数据产业高质量发展的若干措施 (Several Measures on Accelerating the High-quality Development of the Data Industry), providing for coordination with professional institutions, including the Beijing Internet Court and Beijing Lawyers Association, to promote the implementation of liability-exemption mechanisms; see also Regulation (EU) 2024/1689, Art. 57(12). ↩