

SUBMISSION TO THE PERSONAL DATA PROTECTION DEPARTMENT (JPDP)

Public Consultation Paper No. 1/2026: Artificial Intelligence (AI) and Personal Data Protection Framework

Response of Edwin Lee & Partners

Submitted by: Edwin Lee & Partners, Kuala Lumpur **Contact:** Edwin Lee, Founder and Managing Partner, edwin@lpplaw.my **Date:** 22 October 2026 **Reference:** Kertas Konsultasi Awam Bil. 1/2026, Kerangka Kecerdasan Buatan (AI) dan Perlindungan Data Peribadi

Summary of Recommendations

We make eleven recommendations, each developed in the Parts that follow. A reader short of time can take our position from this table and the boxed recommendation that closes each Part.

No.	Subject	Recommendation in brief
1	Status and effect	Add a clause stating the Framework's legal effect (comply or explain) and its relationship to the Commissioner's other instruments
2	Scope and proportionality	Exclude rules-based systems; give "proportionate" a two-level reference point; map the AI Service Provider onto processor and third party
3	ADMP consistency and oversight	Align with the ADMP Guideline on DPIA, DPO involvement and the no-sole-factor rule; add a human-oversight standard
4	Training on existing data	Use the Act's "directly related" test in paras 6.3, 7.4 and 8.3(d); say how it applies to training; coordinate with MyIPO; refer the reform question upward
5	Material change and testing	Give indicative triggers for "material change" and ask that acceptance criteria be set before testing
6	Procurement of external AI	Role determination, consumer-tier terms, unobtainable due diligence, a fuller TIA for offshore services, local deployment of open models, model clauses
7	Data subject rights	Clarify inferences, model weights, logs, correction, withdrawal and portability; refer to section 42; do not imply rights the Act does not confer
8	Fairness, security and breach	Anchor fairness to the Principles; add four AI-specific fields to the breach form; set a processor-to-controller notification window
9	Records and SME support	One AI Processing Record in place of three lists; a published template; a transition period; a sandbox aligned with NAIO
10	Other regimes	State the boundary with the AI Governance Bill; adopt a one-assessment principle; local instruments first among references
11	Biometric data	Name biometric data, sensitive since the 2024 amendments; state that AI can create it; connect the section 40 and DPIA safeguards

Part 1: Who We Are and Why We Are Writing

1.1 Edwin Lee & Partners is a corporate, commercial and technology law firm in Kuala Lumpur. Most of our work is in data protection: the Personal Data Protection Act 2010 (Act 709), the 2024 amendments, the Commissioner's guidelines and circulars issued since 2024 and, more and more, the use of AI where personal data is involved. We thank JPDP for consulting on the Framework before issuing it. This submission answers the seven questions in the Consultation Paper.

1.2 Our interest. We should say where we stand. We advise Malaysian and international organisations on Act 709 compliance. We also act as the named Data Protection Officer for some clients, so we would be the officer the Framework refers to at paragraph 35.4. And we are a data controller ourselves: we use generative AI in our work, we produce AI-generated video and voice content, and we have built compliance tools for clients' own use, including a DPIA questionnaire. The Framework would apply to us directly, and we have read it with that in mind (see paragraphs 3.3 and 6.2).

1.3 Our earlier submissions. This is the third AI consultation we have responded to this year. On 31 July 2026 we filed a submission on the proposed AI Governance Bill with the National AI Office,¹ and on 14 August 2026 a submission on the proposed Copyright Act amendments with MyIPO.² In both we made the same basic point: most AI systems that matter process personal data, Act 709 and the Commissioner's instruments already govern that processing, and any new instrument should connect to them rather than sit beside them. Several of the things we asked NAIIO and MyIPO for were really matters for JPDP to decide. The Framework is JPDP's answer, and we have tried to respond in a way that is consistent with what we said before. Annexure B shows the correspondence.

1.4 What this submission covers. We respond to all ten Parts of the Framework and all seven questions, but not evenly. We go deepest where our day-to-day work gives us something useful to say: the Framework's legal effect, how it fits with the ADMP Guideline, the legal basis for training on data a controller already holds, the buying of external AI by organisations that will never train a model, the records controllers must keep, and one category of data the Framework does not mention. We say less about development and testing, where others know more. We do not argue from the GDPR. Foreign material appears only in answer to question 7, and only for comparison.

1.5 What the Framework gets right. Our overall view is positive. The Framework stays close to Act 709 and does not bring in ideas the Act does not have. Paragraph 5.1 says that using AI creates no separate legal basis, which is correct and worth saying. Paragraphs 3.2 and 4.5 to 4.8 on anonymisation, pseudonymisation and re-identification are clear. Part F on procurement deals with the situation most Malaysian controllers are actually in. Paragraph 22 handles offshore AI processing through section 129 and the Transfer Impact Assessment instead of a new data sovereignty rule, which is what we asked NAIIO to leave to the Commissioner. Paragraph 31 has a fairness obligation built from Malaysian materials, which we asked NAIIO to put in the Bill. We support the Framework. Our recommendations are meant to help it work in practice.

1.6 How the paper is organised. Each Part below sets out the issue, the applicable law, our analysis and the recommendation. Annexure A maps the recommendations to the seven questions. Annexure B lists the provisions and instruments we rely on and shows where this submission matches our earlier ones. Annexure C is a terms review checklist and model clause schedule for the arrangements at paragraph 21.3. Annexure D offers suggested wording for the paragraphs our recommendations touch. Annexures C and D are offered for the drafting team's convenience and can be used or set aside.

Part 2: The Status and Effect of the Framework

Issue

2.1 The Framework is issued under subsection 48(g) of Act 709 as a Guideline (paragraph 2.1). Paragraph 1.7 says it supplements the Act and does not limit or modify any obligation under it, and paragraph 1.8 describes it as high-level best practice. But the operative text is written almost entirely in mandatory terms: "hendaklah" in Bahasa Malaysia, "shall" in English. The Framework does not say what happens if a data controller does not follow it. Question 5 of the Consultation Paper asks what challenges organisations might face "if the requirements in this Framework are enforced under Act 709", but the Framework does not explain how that enforcement would work.

Applicable law

2.2 Subsection 48(g) empowers the Commissioner to monitor and supervise compliance with the Act, "including the issuance of circulars, enforcement notices or any other instruments to any person". A Guideline issued under it is not a code of practice under sections 23 and 24, non-compliance with which is an offence under section 29, nor a Regulation. The Principles themselves are binding under section 5, and contravention is an offence under subsection 5(2), carrying since the 2024 amendments a fine of up to one million ringgit, imprisonment of up to three years, or both.

Analysis

2.3 In our experience, a guideline written in mandatory terms, with no statement of its effect, gets read in two different ways. Larger controllers with lawyers will treat every "shall" as binding and budget for it. Smaller controllers, which are the majority in Malaysia and the majority of our clients, will notice that it is a guideline and treat it as optional. The first group will complain about cost. The second group will not adopt it at all. Neither outcome is good for data subjects.

2.4 The Framework also joins a set of instruments that is now quite large: the Standard 2015, the Regulations, the DPIA, DPbD, ADMP, DPO and Cross-Border Transfer Guidelines, and Circulars No. 1/2025 and No. 2/2025. It refers to some of them (paragraphs 11.3, 22.2, 33.3 and 37.3) but not others, and it does not say which one governs where two overlap. Part 4 describes one overlap where the Framework and the ADMP Guideline now say different things.

2.5 We think the answer, within a guideline, is a short "status and effect" clause saying three things: the Framework sets out the Commissioner's view of what the Principles require when personal data is processed through AI; the Commissioner will take compliance with it into account when assessing compliance with the Principles; and a controller that departs from any provision should be able to explain why in writing. This is a comply-or-explain approach. It gives the Framework real weight, because an unexplained departure becomes evidence that a Principle has been breached, while being honest about what a guideline is. It also gives smaller controllers a proper way to be proportionate without being non-compliant.

Recommendation

Recommendation 1. The Framework should include, in Part A, a clause stating its legal effect: that it records the Commissioner's view of what Act 709 requires where personal data is processed through AI systems; that compliance with it will be taken into account in assessing compliance with the Personal Data Protection Principles; and that a data controller departing from any provision should document its reasons. The same Part should state how the Framework relates to the Commissioner's other instruments, and should be accompanied by a consolidated cross-reference table identifying, for each topic the Framework touches, the instrument that governs.

Part 3: Scope, Definitions and Proportionality

Issue

3.1 Three questions arise under Part B of the Framework. First, whether the definition of "AI system" at paragraph 3.2, read with paragraph 4.1, reaches software that no one would describe as AI. Second, what the Framework's repeated requirement of proportionality is proportionate to. Third, how the Framework's "AI Service Provider" relates to the data processor and third party under Act 709.

Applicable law

3.2 Act 709 regulates data controllers, data processors and third parties, each defined in section 4. Since the 2024 amendments, a data processor is directly subject to the Security Principle. The Act contains no concept of an "AI Service Provider". The Framework's definition at paragraph 3.2 is "any person or organisation other than the data controller" that develops, provides, hosts, operates, maintains or otherwise supports an AI system for the controller's use.

Analysis

3.3 The definition reaches deterministic software. Paragraph 3.2 defines an AI system in terms that follow the OECD formulation: a machine-based system that infers from input how to generate outputs such as predictions, recommendations, content or decisions. Paragraph 4.1 then provides that AI systems "may use software, algorithms, machine learning models, predefined rules or other computational methods". The reference to predefined rules is the problem, and we can show why from our own practice. The DPIA questionnaire tool we have built for clients is 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. It clearly generates recommendations. On the wording of paragraphs 3.2 and 4.1 it is an AI system, and we would have to test, validate, monitor and document it under Parts D, E and J. We do not think that is what JPDP intends. One sentence would fix it: systems that run only on fixed, pre-defined rules, without learning from data or inferring from patterns, are outside the definition. We made the same point to NAIIO about the Bill. The Framework's wording is actually wider than the Bill's here.

3.4 Proportionate to what? The Framework requires measures to be "berkadar" or proportionate at paragraphs 20.3, 30.2, 32.2, 35.2, 37.1, 40.3 and 41.2, and the factors listed (nature, scope, context and purpose of processing; sensitivity; potential risk to data subjects) are the right ones. What is missing is a reference point. A compliance officer whose company has bought a CRM with an AI recommendation feature cannot tell whether Parts D to J apply in full to that feature or whether a short record is enough. We

are not suggesting a risk classification system; that is for the Bill. We suggest a simple two-level reference point, using a line JPDP has already drawn. The ADMP Guideline applies where an automated decision may produce legal effects for the data subject or significantly affect the data subject (paragraph 7.1), and it makes a DPIA mandatory wherever automated decision-making and profiling is present (paragraph 6.1). The Framework could provide that where an AI system makes, or materially supports, decisions with legal or similarly significant effects on data subjects, the Framework applies in full; and that for other AI systems, Parts C, F, H and a proportionate record under Part J apply, with the remaining Parts as good practice. That gives the word "proportionate" something to attach to, and it uses JPDP's own line.

3.5 The AI Service Provider. The Framework's definition cuts across the Act's categories. Paragraph 21.2 correctly asks the controller to determine whether the provider acts as a data processor or as a third party, but gives no guidance on how. In the configurations we see most often, an enterprise generative AI provider is a processor for inference on the customer's data and an independent controller for any use of that data to improve its models, and can be both under one contract. Paragraph 1.4 extends references to the data controller to processors and third parties "where applicable" without saying which obligations are meant; since a processor is subject to the Security Principle only, the Framework should say so. A short table mapping the AI Service Provider onto data processor and third party, with the obligations attaching to each, would resolve both points. The same table should distinguish an external AI system accessed as a hosted service from a model deployed on infrastructure the controller controls, a distinction we develop at paragraph 7.9.

3.6 Anonymised data. The Framework defines anonymisation as irreversible processing after which the data subject can no longer be identified, directly or indirectly, having regard to the risk of re-identification (paragraph 3.2). It does not say what follows, which is that properly anonymised data is no longer personal data under section 4 and is outside the Act. Saying so would help controllers deciding whether they can use anonymised data for model development under paragraph 12.4. It would also help to add that a controller sharing de-identified data should prohibit re-identification in the contract, since paragraph 4.8 notes the risk but leaves it with the recipient.

Recommendation

Recommendation 2. Part B should (a) exclude from the definition of AI system any system operating solely on deterministic, pre-defined rules without learning from data or inferring from patterns; (b) provide a two-level reference point for proportionality: an AI system that makes or materially supports decisions producing legal or similarly significant effects on data subjects attracts the Framework in full, and other AI systems attract Parts C, F and H with a proportionate record under Part J; (c) include a table mapping the AI Service Provider onto data processor and third party under Act 709, with the obligations attaching to each; and (d) state that data anonymised to the standard in paragraph 3.2 falls outside Act 709, recommending that recipients of de-identified data be bound by contract not to attempt re-identification.

Part 4: Consistency with the ADMP Guideline, and Human Oversight

Issue

4.1 The Framework cross-refers to the DPbD Guideline, the Cross-Border Transfer Guidelines and the breach notification Circular. It does not mention the Commissioner's Guideline on Automated Decision-Making and Profiling, which is the Commissioner's existing instrument on AI used in decisions about individuals. On three points the two instruments say different things.

Applicable law

4.2 The ADMP Guideline, issued under the same subsection 48(g), provides that where an organisation embarks on automated decision-making and profiling, the Data Protection Officer "shall" be engaged at the earliest possible stage, beginning with a DPIA (paragraph 1.3), and gives the DPO specific responsibilities in overseeing such systems (paragraph 5.1); that ADMP "is one of the qualitative factors that triggers the requirement to carry out a DPIA, regardless of the nature or extent of its intended use" (paragraph 6.1); and, among the best practices for AI in ADMP, that "AI must not be relied upon as the sole factor when making policies and/or decisions concerning a data subject" (paragraph 10.2.6).

Analysis

4.3 The Framework is softer on all three points, and the difference is not explained. On the DPIA, paragraph 37.3 provides that "where appropriate" the risk assessment shall include or be complemented by a DPIA. On the DPO, paragraph 35.4 provides that where a DPO is appointed or required, the DPO shall be "appropriately involved". On oversight, the Framework contains no human-oversight standard at all: paragraph 31.2(d) asks for a review or escalation mechanism "where appropriate", and paragraph 41.3(g) lists "human oversight measures, where applicable" as something an audit may assess. A controller deploying an AI credit-scoring model and reading only the Framework would conclude that a DPIA is optional and that nothing stops the model from being the sole basis of the decision. Reading the ADMP Guideline, it would conclude the opposite on both. We do not think JPDP intends the Framework to relax the ADMP Guideline. If it does, the Framework should say so. If not, the two should be aligned. The same applies to paragraph 42.3(c), which lists DPIAs among the records to keep "where conducted", and so reinforces the reading that they are optional.

4.4 The alignment is simple and follows from the reference point at paragraph 3.4. Where an AI system makes or materially supports decisions with legal or similarly significant effects, a DPIA should be required rather than optional, the DPO should be involved from the design stage as the ADMP Guideline already requires, and the no-sole-factor rule should apply. Elsewhere the Framework's "where appropriate" wording can stay. One clarification is needed alongside this: the Framework should not be read as requiring an organisation to appoint a DPO just because it uses AI. Whether a DPO must be appointed is governed by the thresholds in Circular No. 2/2025 and the DPO Guidelines. The duty to assess and the duty to appoint are separate questions, and the Framework should keep them separate. The DPIA Guideline also allows a suitable person other than the DPO to lead an assessment (paragraph 6.5), which matters for smaller controllers.

4.5 We would add that the Framework needs a standard for human oversight, and that having a person "in the loop" is not by itself that standard. A person who does not understand what the system is optimising for, who does not have the information to judge its output, or who has no authority to override it, is not really exercising oversight. We suggested to NAO that oversight be defined by capability: understanding,

information and authority. The same definition would work for the Framework, which controllers will be applying long before the Bill is in force. The Framework should also recognise that oversight comes in degrees. Newer systems take a series of actions rather than producing one output for one decision, and oversight can range from approving every step, to approving significant steps, to reviewing after the event. In our view the last of these cannot be squared with paragraph 10.2.6 of the ADMP Guideline where the decision has significant effects.

Recommendation

Recommendation 3. The Framework should be aligned with the ADMP Guideline and should cross-refer to it in Parts G, H and I. In particular: (a) paragraph 37.3 should require a DPIA, rather than permit one, where an AI system makes or materially supports decisions producing legal or similarly significant effects on data subjects; (b) paragraph 35.4 should require the Data Protection Officer's involvement from the design stage in such cases, consistently with paragraphs 1.3 and 5.1 of the ADMP Guideline, while making clear that whether a DPO must be appointed is governed by Circular No. 2/2025 and the DPO Guidelines and not by the use of AI as such; (c) paragraph 42.3(c) should refer to DPIAs "required under the applicable guidelines" rather than "where conducted"; (d) Part H should state the rule in paragraph 10.2.6 of the ADMP Guideline that AI is not to be the sole factor in such decisions; and (e) Part H should define human oversight by capability (the person understands what the system optimises for, has access to the information needed to evaluate its output, and is empowered to depart from it) and recognise that the degree of involvement required rises with the significance of the decision.

Part 5: Training and Fine-Tuning on Personal Data Already Held

Issue

5.1 Paragraph 5.1 of the Framework says, correctly, that using AI creates no separate legal basis for processing personal data. The question every data controller then asks is whether it can use the personal data it already holds, collected to provide its services, to train, fine-tune or improve a model. The Framework touches on this in three places (paragraphs 6.3, 7.4 and 8.3(d)) but does not answer it. The three places also state the test in three different ways, and none of them is the test in the Act.

Applicable law

5.2 The Act's test for reusing personal data is whether the new purpose is directly related to the original. Processing must be necessary for, or directly related to, the purpose for which the data was collected (paragraph 6(3)(b)); disclosure for another purpose requires consent unless that purpose is directly related to the original (section 8, read with section 39); and notice must be given before data is used for a new purpose (paragraph 7(2)(c)(i)). Where the new purpose is not directly related, the controller needs consent (paragraph 6(1)(a)) or one of the six grounds in subsection 6(2): performance of a contract with the data subject, steps towards such a contract, a legal obligation, vital interests, the administration of justice, or functions conferred by law. None is a legitimate-interests or product-improvement ground. Sensitive personal data requires explicit consent under section 40, and the only public-availability concept in the Act, paragraph 40(1)(c), is confined to sensitive data made public by the data subject's own deliberate steps. The research

exemption at paragraph 45(2)(c) applies only where the data is not also processed for another purpose and the results do not identify the data subject.

Analysis

5.3 Three tests where the Act has one. Paragraph 6.3 requires subsequent processing to be "consistent with" the purpose for which the data was collected. Paragraph 7.4 requires a new notice where data is used for a purpose "not previously specified". Paragraph 8.3(d) requires fresh consent where the purpose "materially differs" from the one consented to. The phrase "materially differs" appears nowhere in Act 709, and "consistent with" is a looser standard than the Act's. A controller could read paragraph 7.4 as requiring a new notice and paragraph 8.3(d) as not requiring fresh consent, while the Act itself asks a different question from both. The fix is to use the Act's own words in all three places: the processing must be for the purpose for which the data was collected or a purpose directly related to it, failing which notice under section 7 and, where no subsection 6(2) ground applies, consent under section 6.

5.4 What "directly related" means for training. Once the test is stated, the Framework should say how to apply it, because this is where controllers most need guidance and have the least. The relevant factors come from the Act itself: how closely the training purpose is linked to the purpose notified at collection; what the data subject was told and could reasonably expect; whether sensitive personal data is involved; the likely effect on the data subject, including the risk that the system reproduces the data in its outputs; and the safeguards applied, particularly anonymisation or pseudonymisation before training, which paragraph 12.4 asks controllers to "consider" and should ask them to prefer. The Framework should also say that anonymising personal data is itself processing, and needs a basis under the Act until the data stops being personal data. One worked example would do most of the work. A fraud-detection model retrained on the controller's own transaction data, for the fraud-prevention purpose notified at collection, is ordinarily directly related. A general customer-service assistant trained on correspondence collected to handle those customers' own matters, and then offered to other customers or licensed to other businesses, is ordinarily not, and requires notice and consent.

5.5 Public data and the research exemption are not shortcuts. Paragraph 6.3 is the only place the Framework deals with publicly available data, and it does not say the thing controllers most need to hear: that public availability does not remove the requirements of sections 6 and 7, and is not consent. Collecting personal data from another organisation's website may also engage section 130, which makes it an offence to knowingly or recklessly collect personal data held by a data user without that data user's consent. Not every instance of collection from public sources is a section 130 offence, and the Framework should describe the section accurately rather than broadly, but it should mention it. Paragraph 10.4 lists the research and statistics exemption without saying that commercial model development will rarely qualify: a system built for deployment is processed for another purpose, and a model from which personal data can be extracted is not a result that keeps the data subject unidentified. Each of these points needs only a sentence in Part C.

5.6 Why we raise this a third time. We raised the training question with NAI0, where we said the gap was not the Bill's to close and should be raised with the Commissioner, and with MyIPO, where we said a text and data mining exception would clear the copyright gate and leave the Act 709 gate untouched. We realise we are making the same point a third time, and now to the body best placed to answer it. We do so respectfully. The Framework is the right place for JPDP to state the position, and leaving it unstated is the least helpful outcome for the organisations we advise. We also ask that Part C's treatment of training data be developed in coordination with MyIPO, so that the Framework and any text and data mining exception are

read together and one set of source records serves both regulators, as we asked MyIPO at its Recommendation 7.

5.7 What a guideline cannot do, and what the Commissioner can. A guideline cannot create a legal basis the Act does not contain, and we are not asking JPDP to read one in. On the bigger question, whether Act 709 should have a basis for AI development, we take no position, and we think that is right. It is a question about where the balance between innovation and individual control should sit, and it belongs to a consultation on the Act itself, whether by an exemption order under section 46 or by amendment. The National AI Action Plan 2026-2030 already contemplates coordination on revisions to Act 709 (initiative E14). Our request is only that the Commissioner refer the question upward, and that the Framework not be drafted as though it had been answered.

5.8 What this looks like for a small controller. Consider an online retailer with about twenty-five staff, some fifteen thousand customer records and no in-house lawyer, which wants to fine-tune its website chatbot on two years of customer chat logs. Its privacy notice says it may use data to "improve our services" but does not mention AI. Under the Framework as drafted, the retailer cannot tell whether fine-tuning is a purpose "not previously specified", one that "materially differs", or one that is "directly related". The safe reading is to go back to every customer for consent. The likely outcome is that nobody does the assessment at all. Under the approach above, the retailer has one question to answer, a short list of factors to answer it with, a clear instruction to pseudonymise the logs first, and a record to keep under paragraph 40.4. It may still need a DPIA, since the use of innovative technology is a qualitative factor under the DPIA Guideline and the Guideline illustrates that factor with AI, and the Guideline lets the owner or a suitable manager lead it (paragraph 6.5). But it can reach a defensible answer without paying for outside advice. In our experience that is what decides whether a guideline gets followed.

Recommendation

Recommendation 4. Part C of the Framework should (a) restate paragraphs 6.3, 7.4 and 8.3(d) using the Act's own test, so that each asks whether the processing is for the purpose for which the personal data was collected or a purpose directly related to it, failing which notice under section 7 and, absent a ground under subsection 6(2), consent under section 6 are required; (b) set out, after paragraph 12.4, the factors by which a controller assesses whether training, fine-tuning or improving an AI system is directly related to the original purpose, with a preference for anonymised or pseudonymised data and at least one worked example; (c) state at paragraph 6.3 that the public availability of personal data does not displace sections 6 and 7, and that collection from another data user's sources may engage section 130; (d) state at paragraph 10.4 the extent to which the research exemption applies to model development; (e) be developed, so far as it concerns training data, in coordination with MyIPO so that one set of records serves both regulators; and (f) record that the question whether Act 709 should provide a lawful basis addressed to AI development is one the Commissioner will raise with the Minister, including through initiative E14 of the National AI Action Plan 2026-2030.

Part 6: Material Change, Testing and Validation

Issue

6.1 Parts D and E turn on two expressions that are not defined: "perubahan ketara" or significant change (paragraphs 11.4, 23.2) and "kesan material" or material effect (paragraphs 11.4, 18.1, 18.2). Paragraph 14.2(b) asks whether outputs are "sufficiently reliable" and paragraph 14.5 permits deployment once the controller is "reasonably satisfied", without requiring that the standard be set before testing begins.

Analysis

6.2 We say little here, because it is not our area. Two points. First, controllers will need indicative triggers for material change: a change of model or provider, new categories of personal data, a new purpose, a new population of data subjects, retraining on production data, and integration with another system. A list in paragraph 18.1 would stop the phrase being read as "any update". Second, a test against criteria set after the results are known is not really a test. Paragraph 14 should require the criteria and the version under test to be recorded before testing starts.

Recommendation

Recommendation 5. Paragraph 18.1 should include an indicative list of changes that are material: a change of model or provider, new categories of personal data, a new purpose, a new population of data subjects, retraining on production data, and integration with another system. Paragraph 14 should require that the acceptance criteria and the version of the AI system under test be documented before testing begins, and that results be recorded against them.

Part 7: Procurement of External and General-Purpose AI

Issue

7.1 Most Malaysian data controllers will never train a model. They will buy a CRM with an AI feature, deploy a chatbot built on someone else's model, call a general-purpose model through an API, and find out that their staff are already using consumer AI tools. Part F is the Part that will be used most, and we think it is the strongest. Our recommendations deal with the points where, in our work, controllers most often cannot tell what it requires of them.

Applicable law

7.2 A data controller remains responsible for processing carried out by a data processor on its behalf and must ensure the processor provides sufficient security guarantees (section 9); since the 2024 amendments the processor is itself subject to the Security Principle. Disclosure to a third party for a purpose other than the one notified at collection requires consent (section 8). Transfers outside Malaysia are governed by section 129 as amended, read with the Cross-Border Transfer Guidelines. The Framework addresses these at paragraphs 20 to 23.

Analysis

7.3 Role determination (paragraph 21.2). The Framework asks the controller to decide whether the AI Service Provider is a data processor or a third party. That is the right question, but the controller has to answer it against standard terms it cannot negotiate. Guidance on the common set-ups would help: an enterprise or API service that commits not to use customer inputs for training and to retain them only for the service is, ordinarily, a processor; a service whose terms permit use of inputs to improve the provider's models is, in respect of that use, a third party processing for its own purpose; and a provider may be both under one contract. The Framework could also say what evidence the controller can rely on: the provider's data processing terms, its published retention and training commitments, and its model or system cards under paragraph 21.4.

7.4 Consumer accounts and training defaults. The set-up that causes the most trouble in practice is the consumer account. As we understand the terms commonly offered, the consumer tiers of popular AI services may allow the provider to use what users type in to improve its models unless the user opts out, while business tiers and API access usually exclude that use by default. The position differs between providers and changes over time, so a controller has to check the terms that apply to it. Once the set-up is known, the position under Act 709 is clear. When an employee puts customer data into a consumer account whose terms allow training, the employer is still the controller, the provider is a third party as far as the training use is concerned, and the employer has disclosed personal data to a third party for a purpose that was never notified. Section 8 does not allow that without consent, and section 40 does not allow it at all for sensitive data without explicit consent. The customers whose data went into the prompt never agreed to any of this. An opt-out on the employee's own account only helps if it was actually switched on and actually works. The Framework should set this out in one paragraph, because in our experience most controllers have not thought about it this way.

7.5 Three compliant routes, at three costs. The Framework should then give controllers a way to comply that does not assume an enterprise subscription, which many Malaysian businesses cannot afford. There are three routes. First, no personal data goes into the tool at all: names and identifiers are removed or replaced before anything is typed in, which a written acceptable-use policy can require at no cost. Second, where the provider offers a working opt-out from training, every account used for work has it switched on, checked and recorded, and no sensitive personal data goes in. Third, a business tier or API with no-training terms, which is the only route for sensitive personal data or for processing at volume. Put that way, the Framework is realistic about cost without being silent about the law. We would add two requests. JPDP could publish a short checklist of what to look for in a provider's terms (training use and its default setting, retention, storage location, sub-processors, breach notification). And JPDP could speak to the major providers directly about consumer-tier defaults for Malaysian users, as other data protection regulators in the region have done. Neither costs a controller anything, and the second may do more good than any amount of guidance.

7.6 Shadow AI (paragraph 16.4). Paragraph 16.4 prohibits personal data being entered into an AI system where the controller has not authorised that use. That is right, and it describes the most common AI risk we see: staff using consumer AI tools on personal accounts. The Framework states the rule but not the measures that make it work. Paragraph 36.3 should refer to an acceptable-use policy for AI tools, an approved-tool list, staff awareness directed at this behaviour and, for larger controllers, the data loss prevention controls at paragraph 32.5(a), which the Framework does not currently connect to paragraph 16.4.

7.7 Due diligence that cannot be done (paragraph 20.2(c)). The Framework asks the controller to consider the sources and quality of the provider's datasets or pre-trained models "where information is reasonably

available". For general-purpose models it cannot. We raised the other side of this problem with NAIIO: a foreign developer may be prevented by its home law from producing the documents a Malaysian regulator asks for. The controller buying the service faces the same barrier from the other side. The Framework should say what the controller can rely on instead: the provider's published documentation, independent assurance such as ISO/IEC 42001 certification or SOC 2 reports, contractual warranties that the training data was lawfully sourced, and the controller's own monitoring of outputs. Otherwise auditors will read paragraph 20.2(c) as a requirement no controller can meet.

7.8 Offshore services and the Transfer Impact Assessment (paragraph 22). Almost every general-purpose AI service used in Malaysia processes personal data outside Malaysia when it answers a prompt. Paragraph 22.2 asks for a Transfer Impact Assessment "where appropriate". For an external AI service we think it is always appropriate, and the Framework should say so. It should also say what the assessment has to cover, because the factors in the Cross-Border Transfer Guidelines (paragraph 5.4: similar data subject rights and Principles, similar requirements on processing, DPOs, breach notification and data processors, and an enforcing regulator) do not reach the question a controller using a foreign AI service most needs answered. We suggest three additions: where the personal data is stored and processed, including by sub-processors; the law of the destination country on access by public authorities to data held by service providers; and what independent oversight and redress a data subject in Malaysia would have in respect of such access. These factors apply to every destination by the same standard, and the Framework should not name any country. Since the assessment for a widely used service will be much the same for every Malaysian controller using it, a short template, or published assessments of common destinations, would save each SME doing the same work separately. The paragraph 22.4 record should also serve as the section 129 record.

7.9 Running the model locally. There is another option the Framework should recognise. Several capable models are released under licences that let a controller run them on its own servers or on a cloud region in Malaysia. Run that way, no personal data leaves the controller's environment, there is no transfer under section 129, the model's developer never sees the inputs or outputs, and the question of foreign government access does not arise for the data processed. Paragraphs 1.5 and 1.6 bring open-source models within the Framework but treat them the same way as hosted services. The Framework should distinguish an external AI system used as a hosted service from a model run on infrastructure the controller controls, and should recognise the second as a data-protection-by-design measure under paragraph 11. For a controller worried about both cost and where its data goes, this is often the best answer available. The Framework should also be honest about the trade-off: a controller running a model itself takes on the security, maintenance, testing and monitoring work under Parts D, E and H that a hosted provider would otherwise do, and needs to resource it.

7.10 Processor arrangements (paragraph 21.3). The list at paragraph 21.3(a) to (g) is sound. The difficulty is practical. A Malaysian company subscribing to a global AI service does not negotiate anything. It accepts the provider's standard terms and, usually, a published data processing addendum, and the only things it actually controls are the settings in the admin console (whether inputs are used for training, how long data is kept, which region it is stored in, who can log in) and its own records. For that controller, the "appropriate arrangements" that paragraph 21.3 asks for can only be made up of those three things: the provider's terms and addendum, the configuration chosen, and the record of both. The Framework should say so, so that paragraph 21.3 is not read as requiring a negotiated contract that most controllers cannot obtain. Where the controller does hold the pen, as with a local vendor, a systems integrator or a custom build, published model clauses would give it a benchmark. Annexure C offers both: a terms review checklist keyed to paragraph

21.3 for standard-form services, and a short set of model clauses for negotiated arrangements. JPDP is welcome to use or set aside either.

7.11 One point in the Framework's favour. The Framework attaches its obligations to the controller that procures and uses the system (paragraph 1.4(a)). That is the right point of contact for models trained abroad, as we argued to MyIPO, and it fits the Bill's proposed treatment of the Malaysian deployer as the point of compliance where the developer is offshore. We suggest the Framework say so expressly, in coordination with NAIIO.

Recommendation

Recommendation 6. Part F should (a) at paragraph 21.2, describe the common configurations of AI Service Provider as data processor, third party or both, and the evidence a controller may rely on; (b) state that where a service's terms permit the provider to use inputs to improve its models, the provider is a third party in respect of that use and the controller requires a basis under sections 6, 8 and, for sensitive data, 40 before personal data is entered, and set out the three compliant routes described at paragraph 7.5; (c) at paragraph 36.3, read with paragraph 16.4, refer to acceptable-use policies, approved-tool lists, targeted staff awareness and, where proportionate, data loss prevention controls; (d) at paragraph 20.2(c), state what a controller may rely on where training data information cannot reasonably be obtained; (e) at paragraph 22.2, provide that a Transfer Impact Assessment is expected for external AI services processing personal data outside Malaysia, that it addresses storage and processing location, the destination's law on access by public authorities, and the oversight and redress available to a Malaysian data subject, supported by a template or published assessments of common destinations, with the paragraph 22.4 record serving as the section 129 record; (f) distinguish an external AI system accessed as a hosted service from a model deployed on infrastructure under the controller's control, and recognise the latter as a data-protection-by-design measure; (g) state that, for a service supplied on standard terms, the provider's published terms and data processing addendum, together with the controller's configuration of the service and its record of both, may constitute the arrangements required by paragraph 21.3 where they cover the matters in (a) to (g); publish a checklist of terms to review for standard-form services and non-mandatory model clauses for negotiated arrangements; and engage the major providers on consumer-tier defaults for Malaysian users; and (h) state that the Framework's obligations attach at procurement and use in Malaysia, so that foreign-trained systems are reached through the Malaysian controller, in coordination with the National AI Office.

Part 8: Data Subject Rights in AI Systems

Issue

8.1 Part G is carefully done and we support it, especially paragraph 25.2, which extends the right of access to personal data generated by an AI system where it can reasonably be identified and retrieved. But it leaves open some questions that Data Protection Officers will face on the first access request involving an AI system, and we think JPDP can answer them now.

Applicable law

8.2 The rights of access and correction (sections 30 and 34), the right to withdraw consent (section 38) and the right to data portability (section 43A, inserted by the 2024 amendments) apply to personal data as defined in section 4. A data subject may also, by written notice, require a data controller to cease or not begin processing that is causing or is likely to cause unwarranted substantial damage or distress (section 42), and to cease processing for direct marketing (section 43). Act 709 does not confer a general right to object to automated decision-making or profiling as such, nor a right to require human review of an automated decision.

Analysis

8.3 Six clarifications would be useful. First, that inferences, scores and profiles generated by an AI system about an identifiable individual are personal data, and are within the right of access under paragraph 25.2 where they are retrievable. Second, that model weights and parameters are not personal data, so that access requests do not extend to them. Third, that prompt and output logs will often contain personal data and attract retention and access obligations. Fourth, what "addressing the cause" of an inaccuracy means at paragraph 26.2 where the cause is the model rather than the input data; in our view the Framework should require correction of the stored record and prevention of further reliance, without requiring that the model be retrained. Fifth, that withdrawal of consent under paragraph 27.2 requires cessation of further processing through the AI system and deletion from retrievable stores, but not retraining of a model in which the data is no longer identifiable; the Framework can state the outcome without prescribing the method. Sixth, whether AI-generated inferences fall within the right to data portability. Section 43A refers to the data subject's personal data without distinguishing data the data subject provided from data generated about him. In our view the right is intended for data the data subject supplied, and inferences generated by an AI system fall outside it, but the Framework should say which reading JPDP takes.

8.4 The Framework should also be clear about what Act 709 does and does not provide. Part G, read with paragraph 29, could be taken to suggest a general right to object to AI-driven processing or to insist on human review. The Act gives neither as such. What it does give is the section 42 notice, which a data subject can use against processing through an AI system as against any other processing where it causes unwarranted substantial damage or distress, and the section 43 right against direct marketing. Part G does not mention section 42 at all, and it should, because it is the data subject's real remedy against harmful AI processing under the Act as it stands. Saying what the Act gives, and that it gives no more, protects data subjects from a false expectation as much as it protects controllers. Whether the Act should give those rights is a separate question, which we respectfully suggest the Commissioner raise with the Minister together with Recommendation 4(f). We express no view on it here.

Recommendation

Recommendation 7. Part G should state that (a) inferences, scores and profiles about an identifiable individual are personal data within the right of access where retrievable; (b) model weights and parameters are not personal data; (c) prompt and output logs containing personal data attract retention and access obligations; (d) correction under paragraph 26.2 requires correction of the stored record and prevention of further reliance, not retraining; (e) withdrawal of consent requires cessation of further processing and deletion from retrievable stores, not retraining where the data is no longer identifiable within the model; and (f) whether personal data generated by an AI system falls within the right to data portability under section 43A, which in our view it does not. Part G should also state that the right under section 42 to require cessation of processing causing unwarranted substantial damage or distress is available against processing through AI systems; that Act 709 does not otherwise confer a right to object to automated decision-making or profiling or to require human review; that the Framework does not create such rights; and that the question whether the Act should do so is one for the Minister.

Part 9: Fairness, Security and Breach Notification

Issue

9.1 Part H introduces a fairness obligation (paragraph 31), security measures (paragraph 32) and breach handling for AI systems (paragraph 33). We support all three and have three practical suggestions.

Applicable law

9.2 Act 709 contains no fairness principle in terms. The General Principle (section 6), the Data Integrity Principle (section 11) and the Security Principle (section 9) are the closest statutory anchors. Breach notification is governed by section 12B and Circular No. 1/2025, which prescribes notification to the Commissioner within 72 hours, to affected data subjects within seven days of the notification to the Commissioner where the breach causes or is likely to cause significant harm, staged completion within 30 days, and a prescribed set of particulars.

Analysis

9.3 Fairness. We asked NAIO to add a fairness principle to the AI Governance Bill, built from the Commissioner's own materials, and we are glad to see one in the Framework. The difficulty is that paragraph 31.1 reads as a standalone obligation, and the first question a lawyer will ask is which section it comes from. We suggest it be anchored to the General Principle and the Data Integrity Principle, with the measures at paragraph 31.2 presented as what those Principles require where AI makes or supports decisions with significant effects. The obligation then has a home in the Act, and how much it demands can scale with the reference point at paragraph 3.4.

9.4 Security. Paragraph 32 is sound. It would help to cross-refer to the Personal Data Protection Standard 2015, which already sets baseline security measures, and, for national critical information infrastructure entities, to the Cyber Security Act 2024, so that controllers in those sectors do not end up building a third security standard.

9.5 Breach notification. Paragraph 33 correctly ties AI-related breaches to section 12B and the Circular. Two additions would make it work in practice. First, JPDP owns the notification form. We suggested to NAO that four fields be added for AI-related incidents, and the Framework is the natural place to add them for breaches: the AI system and version involved; whether a DPIA was done and its reference; whether human oversight was designed in and actually exercised; and whether an automated decision about an identifiable person was involved. Keeping the Circular's existing fields and adding these four gives controllers a form they already know how to fill in. Second, paragraph 33.5 asks controllers to make sure providers can supply information in time to meet section 12B, without saying what "in time" means. A controller with 72 hours cannot meet the deadline if its provider also has 72 hours. The Framework should say what notification window it expects from processors; the clauses at Annexure C include one.

Recommendation

Recommendation 8. Part H should (a) anchor the fairness obligation at paragraph 31 expressly to the General Principle and the Data Integrity Principle, with its intensity scaled to the significance of the decisions the AI system makes or supports; (b) cross-refer at paragraph 32 to the Personal Data Protection Standard 2015 and, for relevant entities, the Cyber Security Act 2024; (c) provide for four AI-specific fields in the breach notification under Circular No. 1/2025 where an AI system is involved: the system and version, the DPIA reference, whether human oversight was designed in and exercised, and whether an automated decision about an identifiable person was involved; and (d) state at paragraph 33.5 the Commissioner's expectation of the period within which a data processor or AI Service Provider must notify the controller of a breach, so that the controller can meet section 12B.

Part 10: Records, Documentation and Support for Smaller Controllers

Issue

10.1 The Framework asks controllers to keep records in three places that substantially overlap: paragraph 39.3, paragraph 40.4 and paragraph 42.3. Paragraph 40.5 makes the records producible to the Commissioner on request. For a large controller running its own model this is fair enough. For the retailer at paragraph 5.8, or a trading company with two software tools that happen to have AI features, it is three lists describing the same thing.

Applicable law

10.2 Section 44 of Act 709 already requires a data controller to keep a record of any application, notice, request or other information relating to personal data it processes. The DPIA Guideline requires DPIA records to be kept for at least two years after the processing ends (paragraph 12.1), the Cross-Border Transfer Guidelines require records of each transfer and of the Transfer Impact Assessment (paragraph 16), and Circular No. 1/2025 prescribes the records to be kept of breaches (paragraph 4(7)). The Data Protection Officer's responsibilities under the DPO Guidelines include monitoring compliance across all of these.

Analysis

10.3 We suggest the three lists be consolidated into a single "AI Processing Record", specified once, scaled to the reference point at paragraph 3.4, and published as a template. Two design choices would make it much

more useful at no cost to JPDP. The record should be able to form part of the records the controller already keeps under section 44 and the guidelines above, overseen by the DPO, so that controllers keep one set of records and not two. And the entry for sources of personal data at paragraph 40.4(b) should be specified so that it also serves as the training-data source record that any text and data mining exception under the Copyright Act will require, which we asked MyIPO to align with JPDP. One record that serves the Framework, section 44, the DPIA and Cross-Border Guidelines, the breach Circular and in time MyIPO is the most practical proportionality measure available, and it makes JPDP's supervision easier rather than harder.

10.4 On industry readiness (question 6), four things would turn the Framework into something a small controller can actually operate: the AI Processing Record template; a short Transfer Impact Assessment template for external AI services (paragraph 7.8); a guidance note on role determination (paragraph 7.3); and the terms checklist and model clauses (paragraph 7.10). A dedicated enquiry channel at JPDP for AI-related questions would also help, and would cost little. We would also suggest a transition period. Our experience with the DPO requirements is that smaller controllers did not know the rules applied to them until the deadline had passed. Twelve months from publication, with the first six months clearly guidance-only, would give time for the templates to be published and the first records to be built before the Commissioner expects to see them. Finally, any sandbox under the Framework should be designed together with the National AI Office and should make use of existing national sandbox initiatives where possible, so that a controller is not put through two sandboxes with two sets of conditions. We made the matching request to NAO.

10.5 Records kept under Part J will sooner or later be produced in court, including under section 90A of the Evidence Act 1950. They will carry more weight if they identify the AI system and version that acted, the date, and where the record came from. A small addition to paragraph 40.4 would cover this.

Recommendation

Recommendation 9. Part J should consolidate the records at paragraphs 39.3, 40.4 and 42.3 into a single AI Processing Record, published as a template, scaled to the reference point in Recommendation 2, capable of forming part of the records the controller already keeps under section 44 of Act 709 and the Commissioner's guidelines, identifying the AI system and version concerned, with the entry for sources of personal data specified so that it also serves as the training-data source record under any text and data mining exception. JPDP should publish with the Framework a Transfer Impact Assessment template for external AI services, a guidance note on role determination and model clauses, and should provide an enquiry channel for AI-related questions; provide a transition period of twelve months, the first six guidance-only; and design any sandbox jointly with the National AI Office, using existing national sandbox initiatives where possible.

Part 11: Relationship with Other Regimes, and References

Issue

11.1 The Framework does not address its relationship with the proposed AI Governance Bill, the National Guidelines on AI Governance and Ethics (2024), the Cyber Security Act 2024, the Online Safety Act 2025, or the instruments of sectoral regulators. Question 7 asks which local and international references should inform it.

Analysis

11.2 The Bill is reported to be drafted and on its way to Cabinet, with tabling anticipated between the third quarter of 2026 and the first quarter of 2027.³ JPDP is therefore finalising the Framework at the same time as a statute whose text it may not have seen. We asked NAI0 to design the Bill's Tier 2 assessment so that a DPIA satisfies it in respect of personal data, and vice versa, and to designate the Commissioner as Sectoral Lead for AI systems processing personal data. The Framework can make the complementary commitment: that it governs the personal-data dimension of AI systems, that a DPIA under it is intended to satisfy any corresponding assessment the Bill requires, and that the Commissioner will align the Framework with the Bill on enactment. One assessment per system is the biggest cost saving available to Malaysian controllers, and it only happens if both instruments aim for it.

11.3 References. In answer to question 7, we would put the local instruments first, because what the Framework most needs is to fit with them: the ADMP Guideline, the DPIA and DPbD Guidelines, the DPO Guidelines, the Cross-Border Transfer Guidelines, Circulars No. 1/2025 and No. 2/2025, the Personal Data Protection Standard 2015, and the National Guidelines on AI Governance and Ethics (2024), which already name fairness among their principles. Among foreign data protection regulators, the closest comparators are the Singapore Personal Data Protection Commission's advisory guidelines on AI recommendation and decision systems (2024) and on generative AI (July 2026), which address the lawful basis for training, AI-specific notification, responsibility across the model supply chain, and access and correction against a trained model; the Hong Kong Privacy Commissioner's Model Personal Data Protection Framework for AI (2024) and its checklist on employees' use of generative AI (2025), the closest published treatment of the shadow-AI question; and the UK Information Commissioner's guidance on AI and data protection. ISO/IEC 42001 is the certifiable standard most relevant to due diligence under paragraph 20.2. We cite the EU Artificial Intelligence Act for comparison only, on three points: its exclusion of systems operating on rules defined solely by natural persons (Recital 12), its requirement that testing be conducted against pre-defined metrics and thresholds (Article 9(8)), and its documentation requirements in Annex IV. We do not suggest the Framework adopt the EU model.

Recommendation

Recommendation 10. Part A of the Framework should state that the Framework governs the personal-data dimension of AI systems; that a Data Protection Impact Assessment conducted under it is intended to satisfy any corresponding assessment required under the AI Governance Bill in respect of personal data, and vice versa; and that the Commissioner will align the Framework with the Bill on its enactment. The Framework should cross-refer to the local instruments listed at paragraph 11.3, and may usefully draw, for comparison only, on the regulatory materials of the Singapore and Hong Kong data protection authorities identified there.

Part 12: Biometric Data, Which the Framework Does Not Mention

Issue

12.1 The word "biometric" does not appear in the Framework. Since the 2024 amendments, biometric data is sensitive personal data under Act 709, and AI systems are now one of the main ways it gets created. A controller reading the Framework would not learn either fact.

Applicable law

12.2 Act A1727 inserted into section 4 a definition of biometric data as "any personal data resulting from technical processing relating to the physical, physiological or behavioural characteristics of a person", and added biometric data to the definition of sensitive personal data. Sensitive personal data may be processed only with the data subject's explicit consent or under one of the conditions in subsection 40(1); contravention is an offence under subsection 40(3), carrying a fine of up to two hundred thousand ringgit, imprisonment of up to two years, or both, in addition to the exposure under subsection 5(2). The DPIA Guideline uses facial recognition in two of its worked examples of processing that requires a DPIA: AI-driven facial recognition for authentication, as an example of innovative technology, and in-store facial recognition for a loyalty programme, as an example of systematic monitoring, where the Guideline expressly notes the use of biometric data (paragraph 7.6). The ADMP Guideline mentions biometric data once, at paragraph 8.1, in the context of section 40.

Analysis

12.3 The definition is wider than the foreign material most controllers rely on. The Malaysian definition follows the wording of Article 3(34) of the EU Artificial Intelligence Act rather than the GDPR, and it does not require the data to be used, or usable, to identify the person. On its wording, personal data produced by technical processing of a person's voice, face, gait, typing pattern or facial expression is biometric data whether or not anyone is identified by it. We think that is the better reading, and it has a real consequence: an AI tool that scores a caller's tone, reads a shopper's mood from camera footage, or flags a user by typing rhythm is producing sensitive personal data about an identifiable person. We accept that a narrower reading limited to identification is arguable, and we do not put ours forward as settled. It is because the point is open that the Framework should state the Commissioner's view.

12.4 AI creates biometric data as a by-product. Paragraph 9 of the Framework prompts a controller to ask whether it is collecting sensitive personal data. It does not prompt the question that matters for AI, which is whether the system is creating it. Call analytics, retail cameras, fraud tools and hiring tools all produce biometric data from inputs that are not themselves biometric. The retailer at paragraph 5.8, offered a voice-analytics add-on by its call-centre software vendor, has no reason to think that "calls may be recorded" is not enough, or that explicit consent under section 40 is needed before the feature is switched on. Nothing in the Framework as drafted would tell it so.

12.5 The safeguards already exist but are not connected. The Framework's own provisions on sensitive data (paragraphs 9.1 to 9.4), on separate storage and encryption (paragraphs 32.4(e) and 32.5(d)), and on the sensitivity factor in risk assessment (paragraph 37.2(a)) all apply once the controller recognises the data as sensitive. Nothing in the Framework helps it recognise that. Three additions would: a definition of biometric data in Part B by reference to section 4, with the Commissioner's view on the identification question; a statement in Part C that an AI system may create biometric data from images, audio, video or behavioural signals and that such data is sensitive personal data; and an amendment to paragraph 37.2(a) so that the risk assessment considers biometric data created by the system, not only data collected for it. Getting explicit consent in these settings is genuinely hard, and the Framework should say what it expects: a specific notice, a real choice, and, where the purpose can reasonably be achieved without biometric processing, an alternative for data subjects who decline.

Recommendation

Recommendation 11. The Framework should (a) define biometric data in Part B by reference to section 4 of Act 709 and state the Commissioner's view on whether the definition extends to data not used for identification; (b) state in Part C that an AI system may create biometric data from personal data that is not itself biometric, that such data is sensitive personal data, and that explicit consent or another condition under section 40 is required before the processing begins, with a specific notice and, where reasonably achievable, an alternative for data subjects who decline; (c) amend paragraph 37.2(a) so that the risk assessment considers biometric data created by the AI system; (d) require a DPIA under paragraph 37.3 before an AI system that processes or creates biometric data is deployed; and (e) cross-refer to the ADMP Guideline where such a system makes decisions about data subjects.

Part 13: Closing

13.1 We know this submission asks a lot of a guideline. We have tried to keep every recommendation to something that can be done within one: a clause, a cross-reference, a template, a worked example, or a referral to the Minister. Where we have asked JPDP to state the law plainly on a hard question, it is because the organisations we advise need an answer, and JPDP is the body whose answer they will follow.

13.2 We would be glad to help the Research and Risk Assessment Section further as the Framework is finalised, including by providing the templates and clauses referred to above in editable form.

Yours faithfully,

Edwin Lee

Founder and Managing Partner

Edwin Lee & Partners

edwin@lpplaw.my | +6011-5954 1201 | lpplaw.my

22 October 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.

Sources

1. Edwin Lee & Partners, Written Submission on the Proposed Artificial Intelligence Governance Bill, filed with the National AI Office on 31 July 2026, with Annexures A to C. Available at <https://lpplaw.my/ai-governance-malaysia/>. ↗
2. Edwin Lee & Partners, Submission to the Intellectual Property Corporation of Malaysia on the Proposed Amendments to the Copyright Act 1987: Artificial Intelligence, Copyright and Personal Data, filed 14 August 2026, with Annexures A to C. Available at <https://lpplaw.my/copyright-ai-malaysia/>. ↗
3. The Minister of Digital is reported as stating that the draft Bill has been completed and is being prepared for Cabinet, with tabling hoped for in the third quarter of 2026 or, failing that, the first quarter of 2027 (The Star, 11 September 2026, as reported by W.Media, 15 September 2026). ↗