MAH-KAMARIYAH & PHILIP KOH
ADVOCATES & SOLICITORS

Malaysia's Proposed AI Governance Bill

© Mah-Kamariyah & Philip Koh

No part of this document may be reproduced in any form or by any means without prior written permission of Mah-Kamariyah & Philip Koh. This documentand the information contained herein is providedon an "as is" basis and Mah-Kamariyah & Philip Koh disclaims all warranties, express or implied, including but not limited to any warranty that the use of theinformation herein will not infringe any rights or any implied warranties or is fit for a particular purpose. The information provided in this document is general and may not apply in a specific situation and legal advice should always be sought before taking any action based on the information provided.

MAH-KAMARIYAH & PHILIP KOH
ADVOCATES & SOLICITORS

Malaysia's Proposed AI Governance Bill

A Summary on the National AI Office ('NAIO')'s 2026 Public Consultation

In Brief

- On 10 July 2026, NAIO, an agency under the Ministry of Digital, opened public consultation on the proposed AI Governance Bill ('the Bill'), Malaysia's first statute dedicated to AI. Written feedback closed on 31 July 2026.
- Six (6) areas are covered: scope, governance architecture, governance principles, risk framework, incident reporting and the AI sandbox.
- The approach rests on three (3) pillars: institutional oversight, principles and risk.
- No Bill has been tabled. The Digital Minister has indicated tabling in the Dewan Rakyat in early 2027.

The Current Landscape

AI now screens job applicants, scores loan applications, triages patients and drafts the advice customers read. When it gets something wrong, no single law says who answers for it. A company that uses AI across its finance, HR and legal functions answers to several regulators, each with its own approach.

Against this stands a patchwork: AIGE, voluntary guidelines issued in 2024; the PDPA for personal data; and sector rules written before systems could learn. On AI as such, Malaysian law says nothing.

Our July update covered MyIPO's copyright consultation, which asks what AI may learn from. NAIO's paper asks how AI is built, deployed and held to account. Closing that gap is what the paper addresses and why it matters for Malaysia.

The Missing Pieces

Six (6) shortcomings drive the Bill, one for each area it covers:

- **AI is Undefined.** No Malaysian statute defines AI or names who is responsible for it.
- **Oversight is Fragmented.** Each sector regulator applies its own standards and none coordinates across sectors.
- **Ethics are Voluntary.** AIGE binds no one, so compliance depends on goodwill.
- **Risk has No Yardstick.** Nothing in law separates a scheduling tool from a clinical triage system, and current approaches act only once harm is done.
- **Incidents go Unrecorded.** Sectoral reporting runs in silos, and AI failures and near misses need not be reported as such.
- **Testing has No Home.** No sandbox is built around AI, so systems can reach the public before their risks are understood.

Charting Malaysia's Course

The paper is an exercise in balance, not maximum control. Its stated aims are responsible and trustworthy AI, greater certainty for industry and room to innovate, with lower-risk uses left unburdened.

Two models frame the choice. The European Union's AI Act is a single statute with fixed lists: certain practices are banned outright and named uses, such as recruitment, are high-risk by definition. The United Kingdom has no AI statute; existing regulators apply five cross-sector principles within their own remits. Take a hiring tool that quietly marks down applicants from certain postcodes. In Brussels it is high-risk on day one because recruitment is a listed use. In London it falls to the regulators and equality law already in place.

Malaysia takes a hybrid. Like the EU, it proposes a statute and a central authority; like the UK, it leads with principles and works through sector regulators. Its five principles track the UK's closely, save that human dignity is added and fairness, a standalone UK principle, appears only within data governance. Its own twist is the tier test: a system is unacceptable only where there is intent to cause harm, and high-risk where it creates a risk of one of four defined harms. The postcode tool would be high-risk here only if its unfairness can be tied to death, bodily injury, a constitutional liberty or a breach of written law.

The Proposals in Detail

Each row sets out one change compared against the position today.

Position today	What is Proposed	Why Malaysia needs it
1. Scope of the Bill		
No statute defines AI. AIGE and each sector regulator use their own wording.	Functional definitions of AI, AI System and AI Lifecycle, based on what a system does rather than a closed list of techniques, with the eight lifecycle stages of ISO/IEC 5338:2023 as a reference.	Gives every regulator one vocabulary that will not date as techniques change.
Responsibility for an AI outcome must be argued case by case under contract, tort and sector rules.	Two duty-holders: the Developer, who shapes what the system can do, and the Deployer, who decides where and how it is used, each accountable to the extent of its control at each lifecycle stage.	Tells every Malaysian business which duties are its own, whether it builds, adapts or simply uses AI.
No AI-specific rule reaches a system hosted offshore.	The Bill reaches systems placed on the market or put into service in Malaysia; designed, developed or used in Malaysia; or used by a Deployer established in Malaysia, wherever the system is hosted.	Stops offshore hosting being used to sidestep Malaysian obligations.
Private and commercial use of AI are equally unregulated.	Exempt use by an individual for personal, family or household affairs, and systems used solely for national defence or security.	Keeps the law out of private life. The word "solely" keeps dual-use systems in scope.
2. AI Governance Architecture		
No body holds statutory powers over AI as such.	A Central AI Authority as the principal national body, operationalising baseline principles and standards and closing capacity gaps between sectors.	One baseline for the company whose AI runs its finance, HR and legal functions at once.
AI safety work has no statutory footing.	An AI Safety function: maintain a risk framework, supervise assessments, support testing, develop incident mechanisms, engage in international technical cooperation and research.	Makes safety a public function rather than a matter of corporate goodwill.

Position today	What is Proposed	Why Malaysia needs it
No regulator can investigate an AI failure as such.	An Investigation and Enforcement function: technical fact-finding, governance instruments (subsidiary regulations, codes, standards, guidelines), directions and orders, interim risk-control measures and administrative penalties.	Gives the framework teeth, although the level of penalties is not yet stated.
Guidance on AI compliance is scattered.	An AI Enablement function: guidance, templates, training and practical support for public authorities, regulators and regulated organisations, and the operation of AI sandboxes.	Helps SMEs and public bodies comply without building expertise from scratch.
Each regulator acts alone on AI.	Sectoral Leads: existing regulators with the legal authority, technical expertise and governance capacity, delegated powers to turn the national framework into sector codes and guidelines.	Builds on existing expertise instead of placing a second regulator over every sector.
3. AI Governance Principles		
AIGE principles are voluntary and bind no one.	A statutory duty on Developers and Deployers to have "due regard" to five principles throughout the lifecycle, proportionate to the system's nature, context and impact, phased in through voluntary instruments before full codification.	Turns guidance into a duty while giving industry time to adjust.
No rule requires human review of important AI decisions.	Principle 1 (Human dignity and agency): AI must not remove meaningful human agency or place important decisions entirely beyond human review.	Keeps people, not systems, as the subjects of law, rights and responsibility.
No general duty to tell people that AI shaped a decision about them.	Principle 2 (Transparency and explainability): proportionate to context, risk and those affected; those overseeing AI must understand its role, limits and impact.	Gives affected persons awareness and safeguards where AI has material effects.
Blame can be deflected onto "the system".	Principle 3 (Accountability and redress): responsibility stays with identifiable persons through a traceable chain, including the ability to suspend or withdraw a system.	Ensures a person always answers for an AI outcome.
Testing before release is a commercial choice outside regulated products.	Principle 4 (Safety and robustness): anticipate, assess and address risks before they materialise; systems must withstand misuse, interference or failure.	Moves safety from repair after the event to design before it.
The PDPA covers personal data only; nothing governs where other training data comes from.	Principle 5 (Data governance): data must be fit for purpose, properly sourced, secure and managed in accordance with applicable laws.	Dovetails with MyIPO's consent and transparency proposals on AI training data.
4. AI Risk Framework		
No legal test says when AI is risky.	Four baseline harms: death; bodily injury; unlawful deprivation of a fundamental liberty under the Federal Constitution; and contravention of any written law, expandable for each sector.	Anchors risk to harms Malaysian law already recognises.

Position today	What is Proposed	Why Malaysia needs it
Where risk is assessed at all, it is a one-off label.	Risk measured by likelihood, severity and scale, and duration and reversibility, across the whole lifecycle.	Treats risk as dynamic, so a system is reassessed as its use changes.
All AI is treated alike.	Three tiers. Unacceptable: intent to cause harm, prohibited with immediate intervention. High: no intent but a risk of harm, with risk assessment, documentation, internal controls, human oversight, monitoring and mitigation. Low: baseline due regard.	Concentrates obligations where harm is likely and leaves low-risk innovation free.
Sector rules apply regardless of the technology.	Three levers, chosen by the Authority or the Minister: by Class of AI Systems (the Developer is the primary duty-holder), by Sector (through Sectoral Leads) and by Domain of Deployment (the Deployer is the primary duty-holder). Requirements may be mandatory or voluntary.	Lets the regulator target the technology, the industry or the use case, as the problem demands.
5. AI Incident Reporting		
Sectoral reporting runs in silos with no AI-specific coordination.	An AI incident covers failures, weaknesses, misuse, unexpected effects and near misses, not only harm already done.	Captures early warnings before they become harms.
There is no AI-specific complaint channel.	Developers and Deployers report through prescribed channels; the public complains through the Authority. Reports cover the nature of the incident, foreseeable harm, containment, root cause and remediation.	Catches harms that would never be disclosed voluntarily.
Lessons from AI failures are not pooled.	The Authority builds technical records, oversees confirmation, response and recurrence prevention, issues findings and may keep a central repository. Existing sector mechanisms are used, not displaced.	Turns single incidents into national learning without duplicating existing reporting.
6. AI Sandbox		
No supervised testing environment is designed around AI.	A controlled environment to test a system's full context of use, run centrally by the Authority or through designated sectoral sandboxes.	Gives Malaysian start-ups and SMEs a supervised route to market without the full regulatory weight on day one.
Regulators learn about AI risk only after deployment.	Nine potential functions, from isolated testing and development platforms to controlled data access, model evaluation, limited-user pilots, workflow trials, simulated environments and testing under relaxed rules.	Produces evidence to refine the rules before they apply at scale.

Key Terms and Abbreviations

AI: Artificial Intelligence: a functional capability or set of methods enabling a system to perform functions resembling human cognition, such as learning, reasoning, prediction, language processing, content generation or decision support.

AI Lifecycle: the journey of an AI System from early development until it is withdrawn.

AI System: the operative system, whether standalone, embedded, remotely hosted or delivered through an API, through which AI functionality is used in practice.

AIGE: the National Guidelines on AI Governance and Ethics, voluntary guidelines issued by MOSTI in September 2024.

API: Application Programming Interface.

Central AI Authority: the proposed principal national body for AI governance.

Deployer: the person or organisation that causes an AI System to operate in the real world and decides whether, where, how and under what conditions it is used.

Developer: any person or organisation that materially shapes what an AI System can do, how it functions and performs, its built-in limits or the risks it may create.

ISO/IEC 5338:2023: the international standard on AI system lifecycle processes.

MOSTI: the Ministry of Science, Technology and Innovation.

NAIO: the National AI Office, under the Ministry of Digital.

PDPA: the Personal Data Protection Act 2010.

Sandbox: a controlled environment to test an AI System's purpose, limits, safeguards, human oversight, user interaction and real-world effects.

Sectoral Lead: an existing regulator delegated powers under the Bill for its sector or domain.

SME: Small and Medium Enterprise.

What Happens Next

The reform is structural not incremental. AI gains a statutory home. Businesses gain a clear answer to who is responsible. Regulators gain a common baseline and the powers to enforce it. The public gains a channel to report harm. Innovators gain a supervised place to test.

What remains open is the detail: where the Central AI Authority will sit and how independent it will be; the level of penalties; the thresholds between tiers; which breaches of "any written law" trigger high-risk status; whether incident reports will carry a deadline, as the PDPA's 72-hour breach notice does; and how a participant leaves the sandbox. How "due regard" is read also matters: in English public law the phrase imposes a duty to consider, not a duty to achieve a result. The paper does not make fairness a principle of its own, although it cites bias in hiring as its own example.

The draft Bill is now before Cabinet. Businesses that map their AI use now, and decide whether they act as Developer, Deployer or both, will be ready to engage when the text is tabled.

Key Takeaways

- AI gets a statutory home: one Central AI Authority, many Sectoral Leads.
- Responsibility follows control: Developers and Deployers answer, not the machine.
- Principles come first; detailed rules follow in phases.
- Risk turns on harm and intent, not a fixed list of uses.
- Failures and near misses become reportable.
- It is still only a draft Bill.

This legal update is by Kenneth Chong Kheng Aik (Partner) and Ong Li Jun of Mah-Kamariyah & Philip Koh's Intellectual Property, TMT & Data Protection team. This legal update contains general information only and does not constitute legal advice. For further information please get in touch with a member of MKP's Intellectual Property, TMT & Data Protection team at general@mkp.com.my.