[FEEDBACK] AI Governance Bill - Public Consultation - Asia Tech Alliance (ATA) - 30 July 2026

Executive Summary

The Asia Tech Alliance (ATA) welcomes the opportunity to provide input on the draft AI Governance Bill. ATA supports the government's underlying policy objectives: establishing a coherent, cross-sectoral framework for the safe and trustworthy use of AI in Malaysia, while providing industry, government, and society with greater regulatory certainty in developing, deploying, and adopting AI. Malaysia's role as a growing regional hub for AI development and adoption depends on a governance framework that is clear, proportionate, and adaptable to a fast-evolving technology.

ATA's comments focus on several areas where further clarification, calibration, or industry engagement would help ensure the framework achieves its objectives while streamlining compliance obligations. These include clarifying the scope and definitions applied to AI, AI Systems, Developers, and Deployers, the allocation of powers and responsibilities between the proposed Central AI Authority and Sectoral Leads, the design of the AI incident reporting mechanism, and the considerations for lower-risk applications.

Positive elements to acknowledge

- **Adopting a principle-based, technology-neutral approach**
    - ATA welcomes the proposal to anchor the framework in five broad AI Governance Principles rather than prescriptive technical rules. This allows the primary legislation to remain stable and relevant as AI technology evolves, while enabling the Central AI Authority to develop and refine operational guidance, standards, and codes of practice over time without the need to amend primary legislation.

- **Adopting a proportionate, risk-based framework**
    - ATA appreciates that the Bill does not impose uniform obligations on all AI Systems, but instead calibrates regulatory expectations to the level of risk posed, using a harm-anchored, three-tier framework. This risk-based approach helps ensure that lower-risk AI applications are not unnecessarily

burdened, while higher-risk use cases affecting safety, rights, or public interest receive appropriately stronger oversight.

- **Establishing an AI Sandbox to support innovation**
  - ATA supports the proposal to establish an AI Sandbox, which would allow Developers — including start-ups and small and medium-sized enterprises — to test and refine AI Systems under supervised, controlled conditions before wider deployment. This provides a practical pathway for responsible innovation and can generate valuable evidence to inform the future calibration of the regulatory framework.

- **Committing to a phased implementation approach and further consultation**
  - ATA appreciates that the Bill intends to begin with voluntary instruments to guide implementation before progressing to more detailed or codified requirements as ecosystem maturity develops. ATA also welcomes NAIO's approach of seeking broad stakeholder input at this early stage through detailed guiding questions across each component of the Bill, which reflects a considered, engagement-based approach to policy development.

Key areas for improvement

- **Leveraging and evolving existing sectoral rules to manage AI-specific risks**
  - ATA recognises that NAIO has invited views on the specific powers and functions Sectoral Leads should hold (Guiding Question 5, p.8).
  - ATA notes that many organisations already operate under existing sector-specific obligations — for instance, in financial services, healthcare, or telecommunications — administered by established regulators. AI-related obligations would benefit from clear coordination that improves efficiency in assessments and guidelines development.
  - ATA recommends that, wherever possible, sector-specific AI risks be managed by updating and extending existing sectoral rules and regulatory regimes, rather than developing new, parallel AI-specific requirements administered separately from a sector's established regulator. Many of the risks AI introduces in a given sector — for instance, risks to financial soundness, patient safety, or consumer protection — are extensions of risks

those sectoral regimes already exist to manage. As a matter of regulatory principle, rules of this nature should generally remain technology-neutral: the underlying obligation (e.g., ensuring sound credit decisions, safe clinical outcomes, or fair consumer treatment) should apply regardless of whether AI is involved in the process, updated as needed to address AI-specific considerations. Should sector-specific guidelines and/or codes of practice be issued, we suggest that these guidelines and/or codes of practice be open for public consultation before finalisation.

- ATA also suggests that any expansion of harm categories by the Central AI Authority and/or Sectoral Leads be open for public consultation and accompanied by reasonable transition periods for compliance.

- **Coordinating risk allocation between Developers and Deployers**
  - ATA supports the principle that responsibility for AI Systems should remain attributable to identifiable persons through a clear and traceable chain of accountability (AI Governance Principles, Principle 3, p.9), and welcomes the Bill's existing recognition that obligations should correspond to a party's actual "degree of control" over the AI System at each stage of the Lifecycle (Who is Being Regulated, p.5).
  - ATA notes that Developers and Deployers are best placed to manage different dimensions of risk. Developers are generally better positioned to provide clear, current technical documentation on an AI System's capabilities, limitations, and intended parameters of use — recognising that such documentation evolves as models and systems are updated. Publication of such documentation is already a practice today. Deployers, in turn, are generally better positioned to assess the context-specific risks arising from a given use case, including the affected users, the operational environment, and the safeguards appropriate to that context.
  - ATA recommends that subsidiary instruments allocating specific duties between Developers and Deployers reflect each party's practical ability to manage the relevant risk, industry best practices, and allow such practices to evolve in tandem with AI technological advancements. The speed and level of detail prescribed in the subsidiary instruments would ideally reflect this understanding as well. In this regard, in developing the subsidiary instruments, ATA proposes that Developer obligations should center on

providing accurate, accessible, and up-to-date technical documentation, while Deployer obligations should centre on context-specific risk assessment and safeguards.

- As a guiding principle, ATA recommends that NAIO consider a regulatory focus on the actual outcomes and outputs an AI System produces — including whether outputs disclose personal data, generate harmful content, or facilitate unlawful activity. An outcomes-focused approach allows Developers the flexibility to apply the most effective and current technical methods while Deployers can focus on assessing and governing the real-world outputs and actions their AI Systems produce in context.

- **Calibrating the thresholds for Developer and Deployer status**
  - ATA welcomes the Bill's distinction between Developers and Deployers as a practical basis for establishing clearer accountability across the AI value chain (Who is Being Regulated, p.5). As NAIO develops the framework further, ATA recognises the need to specify clearer thresholds for the scenarios in which organisations are treated as a Developer or Deployer based on the relative allocation of control over the AI System at each stage of the Lifecycle.
  - The Bill currently identifies a party that "adapts [an AI System] for a specific use case" or "later modifies it after deployment" as a Developer, which on its face could capture routine downstream adaptation or task-specific fine-tuning undertaken by an organisation that is, in substance, a Deployer.
  - ATA recommends that subsidiary instruments set an appropriately calibrated threshold for when a party acquires "Developer" status. The threshold should turn on when a modification materially changes the AI System's capabilities or safety profile, rather than on whether the underlying model has been technically modified. Common adaptation techniques necessarily involve some modification without materially altering the system's risk behaviour. This would ensure obligations designed for parties that materially shape an AI System's capabilities are not inadvertently extended to organisations undertaking ordinary contextual customisation.
  - ATA also notes that the "Developer" definition, as currently framed, could equally extend obligations to providers of general-purpose AI models, even where such providers have no visibility into or control over downstream deployment decisions. This also applies to models distributed under

open-source or open-weight terms. The original provider is not in a position to implement real-time post-deployment monitoring or system-level safeguards as these are architecturally implemented by the deployer at the system level.

- ATA recommends that the framework recognise a proportionate distinction between providers of general-purpose models and developers of application-specific AI Systems. Obligations on upstream model providers should be calibrated to what they can reasonably control — such as pre-release safety evaluation, clear documentation of a model's capabilities and limitations, and the licence terms under which it is released — rather than extending system-level deployment obligations to parties that do not determine the context of use.
- Upstream providers should be expected to furnish information sufficient for downstream Deployers to meaningfully assess risks and implement safeguards in their intended context of use. This ensures that calibrating upstream obligations does not leave Deployers, particularly smaller organisations, bearing responsibility for risks they have no practical means of evaluating.
- In the context of open digital marketplaces, the definition of “Deployer” should be clarified for platforms hosting user-generated content. Where a third-party seller independently uses AI tools to generate product listings, images, or descriptions that are subsequently hosted on a marketplace, it is the seller — not the platform — that causes the AI System to operate and determines the conditions of its use.
- ATA therefore recommends for NAIO to clarify that the Bill's existing definition of "Deployer" excludes platforms in this scenario. Absent such clarification, platforms may face uncertainty as to whether they inherit Deployer obligations for AI outputs generated entirely outside their control and visibility. ATA notes that platforms will continue to address unlawful or harmful content through established notice-and-takedown mechanisms and industry provenance standards such as C2PA.

- **Encouraging good-faith incident reporting and clarifying reporting responsibility between Developers and Deployers**

- ATA supports the proposal for an AI Incident Reporting Framework, and appreciates that near misses might provide valuable early warning of system weaknesses even where no harm has yet occurred (AI Incident Reporting, Section 1, p.13).
- ATA encourages NAIO to confirm that the primary purpose of the reporting framework is to support organisational and systemic learning. This would enable the Central AI Authority to identify root causes, recommend mitigations, and strengthen the broader AI ecosystem, rather than to operate as a punitive compliance trigger (AI Incident Reporting, Section 3, p.13). A framework perceived as primarily punishing risks discouraging timely, good-faith disclosure, particularly of near misses, which would undermine the early-warning function identified by NAIO.
- ATA also notes that the paper currently identifies "the Developer or Deployer" as the primary source of incident reports, without further distinguishing which party bears the reporting duty in a given case (AI Incident Reporting, Section 2, p.13). In practice, Developers and Deployers typically have visibility into different parts of an incident. Developers are generally better placed to identify issues originating in the design, training, or underlying behaviour of an AI System. In contrast, Deployers are generally better placed to detect and report issues arising from the operational context in which the system is used, including how it was configured, monitored, or relied upon.
- ATA recommends that subsidiary guidance clarify how reporting responsibility is allocated as between Developers and Deployers, calibrated to which party had visibility into, or control over, the relevant aspect of the incident. This would avoid ambiguity as to which party must report in a given scenario, reduce the risk of duplicate or conflicting reports for the same incident, and ensure reporting obligations remain proportionate to each party's actual role in the AI Lifecycle.
- ATA further recommends that reporting obligations be scoped so that Developers and Deployers are not, in the course of good-faith disclosure, drawn into providing more extensive information than is reasonably necessary to investigate and address the incident — including proprietary model details, source code, or other trade secrets. Clear boundaries around the scope of required disclosure, together with confidentiality safeguards for

information provided, would encourage full and prompt cooperation with the reporting framework while protecting the legitimate commercial interests of reporting parties.

- **Providing transitional arrangements for AI Systems already in deployment**
  - ATA welcomes the lifecycle-based structure of the obligations proposed (Scope of the Bill, Artificial Intelligence Lifecycle, p.4). However, the consultation paper does not currently address how the Bill would apply to AI Systems that are already designed, developed, or deployed in Malaysia at the time the Bill takes effect.
  - Given that many AI Systems currently in production were not built with this framework in mind, applying the full suite of obligations immediately upon commencement could be operationally difficult, particularly for systems with long deployment histories or complex integration into existing business processes.
  - ATA recommends that the Bill include reasonable transitional arrangements for AI Systems already in deployment — such as a grace period to achieve compliance, or a phased application of obligations by risk tier — to allow Developers and Deployers sufficient time to assess and adapt existing systems without disrupting ongoing operations. Aligning requirements in the Bill with international standards, which many organisations already operate in accordance with, would further ease this transition. To ensure a smooth transition for global digital platforms operating unified codebases, the framework should also prioritise international interoperability so that Malaysian consumers retain uninterrupted access to global features.

- **Clarifying scope and providing proportionate exemptions for lower-risk applications**
  - ATA supports the existing exemptions for personal, family, or household use and for national security systems (Scope of the Bill, Section 4, p.5), and notes that NAIO has explicitly invited views on whether further exemptions should apply based on sector, domain, or use case (Scope of the Bill, Guiding Question 3, p.5).
  - ATA recommends that the Bill also consider a proportionate exemption, or a clearly lighter-touch tier, for AI Systems used for genuine research,

educational, or personal non-commercial purposes, provided such use does not present significant risk to individuals or society. This would avoid placing unnecessary compliance burdens on academic institutions, researchers, and hobbyist or early-stage developers, and would support Malaysia's broader innovation and talent-development objectives without compromising the Bill's core safety goals.

- ATA further recommends that implementing guidance distinguish task-bound, transactional AI tools (e.g., conversational search, customer service assistance, and marketplace recommendations) from companion or social AI systems. The former operate within a defined commercial purpose and a bounded interaction, and are not designed for open-ended, sustained relational engagement with users. These present materially different risk profiles and should generally be classified as low-risk, having regard to whether the tool is authorised to take material real-world actions.
- ATA also recommends that internal backend systems (e.g., search ranking, content organisation) and automated tools used for trust and safety purposes (e.g., content moderation, fraud prevention, and intellectual property protection) remain outside the scope of obligations designed for high-risk AI Systems. Over-regulating internal safety mechanisms could inadvertently slow down a platform's ability to swiftly mitigate consumer harm, which would go against the Bill's safety objectives.
- Where such systems do raise risks to individuals, those risks are already addressed by existing data privacy, consumer protection, and sectoral regulatory frameworks. This is consistent with ATA's overarching recommendation that AI-specific risks be managed through existing regimes wherever possible rather than through parallel AI-specific requirements.
- In addition, subsidiary guidelines should clarify that a cross-border platform operating global infrastructure that is merely accessible to Malaysian users is not, on that basis alone, treated as having placed an AI System on the market or put it into service within Malaysia (Scope of the Bill, Section 3, p.5). Establishing clear thresholds for cross-border digital services would provide meaningful certainty for global providers and support continued access to AI services in Malaysia.

- **Providing guidance on enforcement and penalty determination**

- ATA recognises that meaningful administrative penalties are necessary for the Bill's obligations to carry real deterrent effect (The AI Governance Architecture, Investigation and Enforcement Functions, p.8).
- It will be helpful for the industry to understand the Central AI Authority's approach, for instance through published guidance on the factors it will take into account, e.g. the severity and actual impact of the breach, whether it is a first-time or repeated occurrence, the degree of cooperation with the government during investigation, and the timeliness and adequacy of remediation. Enforcement guidance would improve predictability for industry and support more consistent, proportionate outcomes.
- We also suggest that any enforcement powers exercised by the Central AI Authority be accompanied by appropriate process safeguards, including transparent decision-making and independent appeal mechanisms.

About the Asia Tech Alliance

The Asia Tech Alliance (ATA) is a regional industry association formed to amplify Asia-Pacific's voice in the global digital economy and champion technology that reflects the region's values, innovation, and potential. Founded by 12 leading global technology companies, the ATA serves as a trusted platform for governments, industry, and communities to co-create human-centric, innovative solutions to the region's most pressing opportunities and challenges. Through policy engagement, collaborative initiatives, and strategic partnerships, the Alliance is committed to shaping an inclusive, agile, and forward-looking digital future for Asia-Pacific.