

Submission on Malaysia's proposed Artificial Intelligence (AI) Governance Bill

A purpose-anchored runtime governance duty
for high-impact AI



Submission on Malaysia's proposed Artificial Intelligence (AI) Governance Bill:

A purpose-anchored runtime governance duty for high-impact AI

EXECUTIVE SUMMARY

Trustethica supports a proportionate, risk-based AI governance bill for Malaysia. The bill should not stop short of only regulating the model provider, AI system, or assessment standards. Instead, organisations deploying high-risk AI must be duty bound to prove that the live use of AI remains within the purpose, authority, and risk boundary it was approved for.

Principles, impact assessments, inventories and testing are necessary but we must go further. Malaysia can distinguish its framework by making operational evidence the de-facto standard. This means every organisation must be able to answer the operational question at the moment AI produces a consequential output or initiates an action: “is this use authorised now, and can we evidence the decision?”.

CORE
RECOMMENDATION

Require each high-impact AI use to have a versioned Authorised Use baseline and proportionate runtime controls. For agentic or directly consequential systems, the control should operate before external effect where reasonably practicable. Material decisions, exceptions and changes should generate a traceable evidence record.

Minimum standard operating model

01	AUTHORISED USE Purpose, owner, scope, data, model, risk and evidence baseline
02	RUNTIME DECISION Evaluate the live event against the current baseline
03	CONTROL OUTCOME Permit, observe, escalate, deny or suspend
04	OPERATIONAL EVIDENCE Preserve the rule, decision, owner action and result

The policy gaps the Bill must address

SAME MODEL, DIFFERENT RISK Risk changes with purpose, affected persons, data, autonomy and downstream consequence.	APPROVAL DECAYS A point-in-time approval does not prove that users, data, models, tools or workflows remain within scope.
DECLARATION IS NOT AUTHORITY The AI system can describe what it proposes to do but the organisation must hold the authority baseline.	A LOG CAN STILL BE TOO LATE A runtime record created after external effect is evidence of what happened, not a control over whether it should happen.

Five provisions that would make the Bill operational

A risk-based bill should be technologically neutral but operationally precise. The following duties can be expressed at statutory level and implemented through approved codes, standards and sector rules.

- 01 Regulate the AI use, not the model or agent in isolation.

High-impact status should turn on the actual or reasonably foreseeable use including purpose, affected persons, data, degree of autonomy, reversibility, materiality and downstream consequence. A general-purpose model is not inherently high-impact in every context. Provider duties should cover capability and product information while deployer duties should cover contextual use and control.
- 02 Create a statutory Authorised Use baseline.

Before a high-impact use goes live, the deploying organisation should record its authorised and prohibited purposes, accountable business, technical and risk owners, data classes and sources, model, tool and integration, affected decisions, risk limits, human oversight, evidence duties, and material-change triggers. The baseline should be institution-owned, versioned and attestable.
- 03 Require proportionate runtime governance.

The organisation should evaluate material runtime activity against the current baseline. Where AI initiates an action or produces an output that directly affects rights, access, safety, financial position or a regulated workflow, the check should occur before external effect where reasonably practicable. A control outcome must be enforceable, monitoring alone should not be represented as control.
- 04 Make human oversight operational.

The Bill or approved codes should require defined reviewer authority, evidence, response windows, capacity and fallback outcomes. A notification sent after execution, an unmanageable queue or a reviewer without decision rights is not meaningful human oversight. Unresolved high-impact events should default to the organisation's defined safe state.

- 05 Require evidence by construction and regulatory interoperability.
- Organisations should preserve a proportionate, tamper-evident record of the applicable baseline, material runtime decision, rule basis, escalation, owner action and outcome. The national framework should complement Malaysia's Personal Data Protection Act 2010, Cyber Security Act 2024, consumer, employment and sector regimes; national coordination should prevent duplicate reporting while allowing sector regulators to set domain-specific thresholds.

Scope discipline

THE BILL SHOULD AVOID	THE BILL SHOULD ENABLE
• blanket treatment of all generative AI as high-impact;	• proportionate duties based on use and consequence;
• a central registry containing sensitive technical or business detail;	• approved codes and technical standards that can evolve;
• duplicate reporting to national and sector regulators;	• privacy-preserving and customer-controlled evidence;
• prescriptive dependence on a particular model, vendor or architecture; and	• regulatory sandboxes and supervised production pilots; and
• “human in the loop” language without decision authority and timing.	• cross-border interoperability, particularly within ASEAN.

Suggested legislative formulation

The following formulation captures the control objective without hard-coding a technical architecture ensuring the bill remains outcome-focused.

1. Authorised use. An organisation deploying a high-impact AI system must establish, maintain and periodically attest to a documented baseline describing the purpose, scope, accountable ownership, data, system components, risk limits, human oversight and evidence requirements under which the use is authorised.

2. Operational control. The deploying organisation must implement measures proportionate to the nature and impact of the use to determine whether material AI activity remains within that baseline. Where an AI system initiates or directly determines a consequential action, the measures must operate before external effect where reasonably practicable.

3. Change and evidence. A material change to purpose, users, data, model, tools, autonomy, workflow or downstream consequence must trigger reassessment before continued high-impact use. The deploying organisation must retain evidence sufficient to demonstrate the baseline, decision, exception, human intervention and outcome.

Innovation and accountability

Malaysia should not have to choose between innovation and accountability. Precise operating duties make responsible deployment easier. Organisations know what they must control, regulators receive testable evidence, and low-impact uses remain proportionate. The Bill can lead ASEAN by requiring high-impact AI to remain within authorised purpose, not only at approval, but in operation.

90-DAY PILOT

Trustethica is ready to deploy a 90-day purpose-anchored runtime governance pilot in a regulated Malaysian environment. The objective is to establish what is technically enforceable, operationally proportionate and suitable for sector guidance. For further information, visit www.trustethica.com or contact hello@trustethica.com.

DISCLAIMER

This submission is provided for policy and technical discussion only and does not constitute legal, regulatory or professional advice. While reasonable care has been taken in the preparation of this document, no representations or warranties are made as to the accuracy, completeness, or suitability of the information for any specific purpose. References to third-party frameworks, standards, or publications are provided for context only and do not imply endorsement, affiliation, or alignment.