

REPORT_DRAFT.md — Asia AI Governance Compliance-Gap Checker

Project: Asia AI Governance Compliance-Gap Checker
Track: Apart Research Asia AI Safety Hackathon — Governance & Geopolitics
Date: June 2026
Status: Research prototype — not legal advice

⚠ Disclaimer: This is a hackathon research prototype. All compliance assessments produced by the tool are illustrative and informational only. They are NOT authoritative legal advice. The governance matrix is a snapshot of regulations as of June 2026; laws evolve. Always consult qualified legal counsel for actual compliance decisions.

1. Problem & Motivation

The Cross-Border Governance Fragmentation Problem

Asia is experiencing a wave of new AI regulation. Three of the world's four largest economies by population — India (1.4B), China (1.4B), and Vietnam (100M) — have enacted or published significant AI governance instruments in 2023–2026. The ASEAN regional body has issued voluntary guidance covering a combined 680 million people across 10 countries.

Yet these frameworks are deeply fragmented:

- **Different binding status:** China and Vietnam have mandatory laws; India's AI-specific guidance is voluntary; ASEAN's guide is non-binding.
- **Different structural approaches:** Vietnam uses a formal three-tier risk system modelled on the EU AI Act; China uses sector-specific instruments with no unified risk classification; India uses a social-context approach with no pre-deployment gate; ASEAN articulates aspirational principles.
- **Different enforcement logics:** Vietnam and China require pre-deployment assessments; India explicitly rejects this model; ASEAN has no enforcement at all.

For a company deploying an AI product across even two of these jurisdictions, compliance is not additive — it is multiplicative and sometimes contradictory. There is no compliance "passport" that satisfies all four simultaneously.

Why This Matters for AI Safety

Regulatory fragmentation creates three compounding AI safety risks:

1. **Compliance arbitrage:** Products may be deliberately routed through the least-regulated jurisdiction to avoid safety checks, even when serving users in stricter jurisdictions.
2. **Knowledge gaps:** Companies lack structured tools to understand where their product sits in this mosaic of regulation — creating inadvertent non-compliance.
3. **Conflict blindspots:** Some regulatory requirements in one jurisdiction may be structurally incompatible with another's (e.g., China's data localization vs. India's cross-border transfer regime), forcing companies to make implicit choices with safety implications that are rarely made explicit.

This project addresses the second and third problems by building a structured, research-backed comparison tool.

2. Method

Research Approach (Phase 1)

- Conducted targeted web searches for each framework using reputable secondary sources (law firm client alerts from Duane Morris, Baker McKenzie, Hogan Lovells, Kim & Chang, OneAsia Legal; government primary sources from PIB India, ASEAN.org; academic sources from Stanford DigiChina; business intelligence from China Briefing; regional law firms Pertama Partners, Rajah Tann Asia, Drew & Napier).
- Every claim tracked with source URL and access date (2026-06-20).
- Unverifiable claims explicitly flagged as UNCONFIRMED rather than guessed (9 flagged items).
- Research notes reviewed and corrected against more-specific Claude review incorporating primary-source corrections (particularly for India's framework — date corrections, ex-ante vs. ex-post distinction, six pillars, three new governance bodies).

Matrix Structure (Phase 2)

The governance matrix (`governance_matrix.json`) organizes all four frameworks along 12 compliance-relevant dimensions:

1. Applicability (who/what is covered)
2. Binding status (mandatory vs. voluntary)
3. Effective date and grace periods
4. Risk classification system
5. High-risk and medium-risk triggers
6. Transparency/disclosure obligations
7. Data localization requirements
8. Human oversight requirements
9. Pre-deployment assessment requirements
10. Enforcement mechanism and oversight body
11. Penalties
12. Data governance obligations

Every field carries a `source`, `source_date`, and `confidence` tag (`confirmed`, `partial`, `unconfirmed`).

Conflict Analysis (Phase 3)

Seven cross-jurisdictional conflicts were identified and documented in `CONFLICTS.md`. Conflicts were required to meet a conservative bar: they must represent either (a) a legal requirement in one jurisdiction that is difficult or impossible to satisfy simultaneously with another's requirement, or (b) a mandate in one jurisdiction where another has an explicit gap, creating real product-level ambiguity. Surface-level philosophical differences were noted but not counted as conflicts.

Tool Design (Phase 4)

The compliance checker is a single-page HTML/JS web app requiring no server or build system. It uses the following architecture:

- The governance matrix (from Phase 2) and conflict register (from Phase 3) are embedded directly in `checker.js`.
- User inputs a free-text AI product description.
- The tool constructs a grounding prompt: the full matrix + conflict register + product description + structured output instructions.
- The prompt is sent to the **Gemini 1.5 Flash API** (user-provided key, processed only client-side).
- The LLM reasons over the matrix to produce per-jurisdiction status assessments and flag applicable conflicts.
- Results rendered as per-jurisdiction cards with status badges, key obligations, gaps, and conflict details.

Critical design choice: The LLM is explicitly instructed to reason only from the injected matrix — not from its own training knowledge of these regulations. This prevents model hallucination from contaminating the analysis. The matrix is the sole source of truth; the LLM's role is structured reasoning, not recall.

3. Key Findings from Conflict Analysis

Seven genuine cross-jurisdictional conflicts were documented. The four most significant:

Finding 1: The Pre-Deployment Gate Divide (C2)

The deepest structural conflict in Asian AI governance is between Vietnam/China (mandatory pre-deployment assessment gates) and India (explicit rejection of pre-deployment gates in favour of ex-post accountability). This is not a matter of different standards — it is a different theory of how AI safety should be achieved. A product optimised for India's fast-launch model may not legally deploy in Vietnam without a conformity assessment; a product built to satisfy Vietnam first will be over-engineered relative to India's requirements but will not conflict with them.

Implication: Companies with India-first development timelines need Vietnam compliance planning that runs in parallel, not sequentially.

Finding 2: Data Infrastructure Cannot Be Retrofitted (C1)

China's data localization requirement (personal data + "important data" must stay in China on local servers) is an infrastructure requirement, not a policy choice. A product that launches globally with a unified data architecture, then tries to enter China, faces a mandatory rearchitecting — not a policy update. This makes China entry a strategic decision that must be made at architecture time, not after product-market fit is established.

Implication: The data localization conflict with India (which permits cross-border flows) cannot be resolved by policy; it requires physical infrastructure bifurcation.

Finding 3: China's Content Requirement is Structurally Incompatible with Global Policies (C5)

China's requirement that GenAI outputs align with "socialist core values" cannot be generalised to a global content policy. Content that is legal and acceptable — or even constitutionally protected — in India, Vietnam, or ASEAN member states may be prohibited in China. A news summarisation AI, a political Q&A chatbot, or any GenAI product touching governance or current events faces this conflict in an irresolvable form without China-specific content governance.

Implication: No global content policy can simultaneously satisfy China and maintain freedom of expression standards applicable elsewhere.

Finding 4: Same Product, Different Risk Classification (C3)

A hiring/recruitment AI would likely be classified High-Risk in Vietnam (affects lawful rights and legitimate interests), triggering mandatory conformity assessment, human oversight requirements, and National AI Database registration. The same product in India has no national high-risk classification and can be deployed without any pre-launch compliance gate (subject only to DPDP data obligations). This asymmetry is not visible from the product's design — it emerges from the interaction between the product's use case and each jurisdiction's risk framework.

4. Tool Demo — Example Outputs

The tool was tested against four pre-built example product descriptions. Expected outputs based on the matrix:

Example 1: GenAI Banking Chatbot (Financial Advice)

Jurisdiction	Expected Status	Key Drivers
🇻🇳 Vietnam	Likely Non-Compliant (without assessment)	High-Risk classification (financial rights). Mandatory conformity assessment required. 18-month grace period for finance applies until Sep 2027.
🇮🇳 India	Requires Additional Obligations	RBI financial AI rules apply. DPDP Act applies to personal data processing. No pre-deployment gate, but strong data governance obligations.

Jurisdiction 🇨🇳 China	Expected Status Likely Non-Compliant (without compliance)	Key Drivers Measures: mandatory security assessment + algorithm filing + content alignment + content moderation + AI output labeling. Financial sector: dual compliance burden (GenAI + PBOC/CBIRC).
🇦🇸 ASEAN	Voluntary Framework	Transparency, Fairness, Human-centricity, and Accountability principles most relevant. No binding obligations.
Conflicts flagged: C1 (data localization if China+India), C2 (pre-deployment gate), C4 (content labeling), C5 (socialist values).		

Example 2: Hiring/Recruitment AI

Jurisdiction	Expected Status	Key Drivers
🇻🇳 Vietnam	Requires Additional Obligations / Non-Compliant	High-Risk (affects employment rights). Conformity assessment, human oversight, registration required.
🇮🇳 India	Requires Additional Obligations	No high-risk classification. DPDP Act: consent for CV data processing. Caste-bias risk triggers social-context risk attention. Graded liability: deployer accountable for discriminatory outcomes.
🇨🇳 China	Requires Additional Obligations	If using algorithmic ranking: Algorithmic Recommendation Provisions apply. If employer-facing SaaS: PIPL data obligations. Content: must align with mainstream values in recommendations.
🇦🇸 ASEAN	Voluntary Framework	Fairness and Equity, Accountability most relevant.
Conflicts flagged: C3 (risk classification asymmetry – High-Risk Vietnam vs. no classification India).		

Example 3: Deepfake Detection Tool

Jurisdiction	Expected Status	Key Drivers
🇻🇳 Vietnam	Likely Compliant (with obligations)	If the tool itself doesn't generate synthetic media, likely Medium-Risk at most. Notification to MoST.
🇮🇳 India	Likely Compliant	No pre-deployment gate. DPDP applies to media submitted. IT Intermediary Rules favour provenance/detection tools.
🇨🇳 China	Requires Additional Obligations	Deep Synthesis Provisions cover detection tools' obligations. If the platform hosts synthetic media for detection: security assessment, labeling compliance. If purely analytical: lighter burden.
🇦🇸 ASEAN	Voluntary Framework	Security and Safety principle strongly relevant. No binding obligations.
Conflicts flagged: C4 (labeling requirements if tool handles synthetic content).		

Example 4: Facial Recognition Surveillance

Jurisdiction	Expected Status	Key Drivers
	Likely Non-Compliant	Almost certainly High-Risk (biometric data; national security context; affects public safety/rights). Mandatory conformity assessment, human oversight, National AI Database

🇻🇳 Vietnam Jurisdiction	(without full Expected compliance) Status	registration, MoST notification. Key Drivers
🇮🇳 India	Likely Non-Compliant (without addressing gaps)	No national high-risk designation, but: biometric data processing requires explicit consent under DPDP Act; mass surveillance of public without consent likely violates DPDP Act and may attract Data Protection Board action. Constitutional right to privacy (Puttaswamy) creates further legal exposure.
🇨🇳 China	Requires Additional Obligations	Public Security Bureau scrutiny. Biometric data = sensitive personal information under PIPL → heightened consent requirements or government authorization. Data localization mandatory for biometric data.
🇦🇸 ASEAN	Voluntary Framework	Security and Safety, Privacy and Data Governance, Human-centricity principles all implicated.
Conflicts flagged: C1 (data localization), C2 (pre-deployment assessment asymmetry).		

5. Limitations

- Matrix is a snapshot.** AI regulations are evolving rapidly. Vietnam's high-risk system list (Prime Ministerial designation) was not fully finalized as of the research date (June 2026). China's enforcement posture evolves through campaigns. India's new governance bodies (AIGG, TPEC, AISI) are being established with unknown operational timelines. The matrix will be outdated within months.
- Matrix relies on secondary sources.** Due to time constraints, the matrix was compiled primarily from law firm client alerts, regulatory trackers, and reputable legal blogs — not from primary text of all statutes. 9 fields remain UNCONFIRMED and should be verified against primary texts before any compliance decision is made.
- The tool assesses applicability, not compliance depth.** The LLM can identify which obligations are triggered, but cannot assess whether a specific product's implementation actually meets them. "Likely Compliant" means the product description suggests no obvious violations — not that a legal audit would confirm full compliance.
- LLM reasoning is not deterministic.** Different API calls for the same product may produce slightly different assessments. The matrix is deterministic; the LLM reasoning over it is probabilistic.
- Scope is limited to four frameworks.** Many ASEAN member states (Singapore, Malaysia, Indonesia, Thailand, Philippines) have national AI policies not covered in this matrix. The tool does not cover the EU AI Act, US executive orders, or other global frameworks.
- No legal review.** The matrix and conflict analysis have not been reviewed by qualified lawyers in any of the four jurisdictions. Do not rely on any output for actual compliance decisions.

6. Implications

For Companies Operating Across These Jurisdictions

- Treat data infrastructure as a compliance decision.** China's localization requirement must be decided before architecture is set. Once a product launches globally with a unified data infrastructure, China entry requires a rebuild.
- Plan parallel compliance tracks.** Vietnam and China require pre-deployment assessments; India does not. Running India and Vietnam/China launches in parallel requires compliance work to start at the same time as product development, not after.
- Design to the most restrictive standard, with exceptions.** A product that satisfies Vietnam's High-Risk requirements (conformity assessment, human oversight, documentation) will generally satisfy India's and ASEAN's standards — except for China's specific content governance requirements, which cannot be generalized.
- Content governance must be jurisdiction-specific for GenAI.** No single global content policy satisfies China's socialist core values requirement and maintains acceptable expression standards elsewhere.

For Regulators

- Harmonization is not inevitable.** India's explicit rejection of pre-deployment gates represents a principled regulatory divergence from Vietnam and China — not a gap in the framework. Convergence on a single Asian AI governance model is unlikely in the near term.
- The ASEAN Guide's voluntary nature creates compliance gaps.** Until DEFA (anticipated end-2026) formalizes binding interoperability rules, ASEAN-level compliance provides no actual legal protection in any member state.
- Mutual recognition agreements could reduce fragmentation.** If Vietnam and China's conformity assessment results could be mutually recognized, cross-border compliance costs would fall substantially. This is a tractable harmonization target.

For AI Safety Research

The cross-border compliance gap is itself an AI safety issue: products that cannot feasibly comply with all relevant jurisdictions' requirements face pressure to choose the most permissive jurisdiction, potentially creating a race-to-the-bottom dynamic. Documenting these conflicts — as this project does — is a prerequisite for designing multi-jurisdictional safety governance that actually works.