

The Board's Guide to Singapore AI Governance

IMDA's Model AI Governance Frameworks, AI Verify and the PDPA — what every board must know

Singapore has taken a deliberately pro-innovation path on AI: a family of voluntary, principles-based frameworks that signal how AI should be governed, sitting on top of one binding law — the Personal Data Protection Act. For a board, the task is to know which is which, and to govern to the spirit of the frameworks before they harden into expectation. This guide is the plain-English orientation; the companion Singapore AI Governance Toolkit turns it into a working board assessment.

1. The lay of the land

Two things sit under every AI decision in Singapore:

- Voluntary frameworks — IMDA's Model AI Governance Frameworks and the AI Verify testing framework. Not law, but fast becoming the benchmark regulators, customers and partners expect you to meet.
- Binding law — the PDPA governs any use of personal data in AI, with real enforcement and breach-notification duties. Where AI touches personal data, this is not optional.

There is no omnibus AI statute. The board's job is to govern AI to the standard of the frameworks, and to comply with the PDPA where personal data is involved.

2. The three-tier Model AI Governance Framework

Singapore now runs a three-tier family, each layer adding to — not replacing — the last.

| Framework | Issued | What it adds |
|--|--|--|
| Model AI Governance Framework (traditional AI) | 2nd ed. 2020 (PDPC + IMDA) | The base layer. Two principles — decisions should be explainable, transparent and fair; AI should be human-centric. Four areas: internal governance; the level of human involvement (in / over / out-of-the-loop, chosen by a harm matrix); operations management (data and model); and stakeholder communication. |
| Model AI Governance Framework for Generative AI | May 2024 (IMDA + AI Verify Foundation) | Nine dimensions for generative AI: Accountability, Data, Trusted Development & Deployment, Incident Reporting, Testing & Assurance, Security, Content Provenance, Safety & Alignment R&D, and AI for Public Good. |
| Model AI Governance Framework for Agentic AI | Jan 2026 (IMDA) — a world-first | For AI that takes actions, not just gives answers. The risk shifts from wrong answers to wrong actions. Four moves: assess and bound the risks upfront (least-privilege autonomy and access); make humans meaningfully accountable (approval checkpoints); technical controls across the agent lifecycle; and end-user responsibility. |

Board takeaway. If you use generative or agentic AI, the newer frameworks apply on top of the base one. The Agentic framework in particular is a signal: boards are expected to bound what autonomous agents can do before deploying them.

3. AI Verify — the testing framework

AI Verify (from the IMDA-owned AI Verify Foundation) is a testing framework and software toolkit that validates an AI system against eleven internationally-aligned principles: transparency, explainability, repeatability/reproducibility, safety, security, robustness, fairness, data governance, accountability, human agency & oversight, and inclusive growth / societal & environmental well-being.

It works two ways: technical tests (quantitative — fairness, explainability and robustness, run by the toolkit) and process checks (documentary verification of governance processes). A companion, Project Moonshot, benchmarks and red-teams large language models.

Crucially: AI Verify does not certify your AI, define ethical standards, or guarantee a system is safe or bias-free. It validates the claims you make and the tests you can pass — useful evidence for a board, not a seal of approval.

4. The PDPA — the part that is actually law

Where AI uses personal data, the PDPA applies. The PDPC's Advisory Guidelines on the Use of Personal Data in AI (March 2024) set out how.

Training AI. Use personal data with meaningful consent, or rely on an exception. The Business Improvement Exception allows use of data you already hold, without fresh consent, where the purpose cannot reasonably be achieved without identifiable data and a reasonable person would consider it appropriate. The Research Exception is broader (including joint research between unrelated firms) but requires clear public benefit, that results are not used for decisions affecting the individual, and that any publication is de-identified.

Deploying AI. You need a consent or legitimate-interests basis; you must notify individuals (purpose, data used, and the role it plays — model/system cards and layered notices are endorsed); and you must be accountable (documented policies, provided on request; DPIAs and AI Verify are encouraged). Where you can anonymise, do — anonymised data falls outside the PDPA.

If a breach happens. A data breach is notifiable if it is likely to cause significant harm, or affects 500 or more individuals. Assess notifiability without undue delay (30 days is the outer marker), then notify the PDPC within 3 calendar days of determining it is notifiable; notify affected individuals at the same time or after — never before.

5. Financial institutions — the MAS overlay

If you are MAS-regulated, add the financial-sector layer. The FEAT principles (Fairness, Ethics, Accountability, Transparency, 2018) set the tone; the Veritas initiative and toolkit operationalise them; and MAS's Information Paper on AI Model Risk Management (Dec 2024) sets out good practices in oversight, risk systems and model development. These are supervisory expectations, not just guidance — and formal AI risk-management guidelines have been signalled, so confirm the current position with your compliance team.

6. Ten questions every Singapore board should ask

1. Do we have an AI policy and a stated AI risk appetite — signed off by the board?
2. Is there a single inventory of where we use AI, who owns each use case, and its impact?
3. For consequential decisions, what is our level of human oversight — and did we choose it deliberately?
4. Where AI uses personal data, is the PDPA basis (consent or a valid exception) documented?
5. Have our AI systems been tested for bias, robustness and explainability — and by whom?
6. Is generative-AI content labelled, and are outputs reviewed before they reach customers?
7. If we use AI agents that take actions, what can they do, and where does a human have to approve?
8. Do we have an AI incident process — detection, response, and PDPC notification if data is involved?
9. If we are MAS-regulated, do we meet FEAT and the AI Model Risk Management expectations?
10. When did we last review all of this — and what would move the posture from red to green?

Turn these questions into a board decision

The **Singapore AI Governance Toolkit** scores your organisation against these frameworks — an interactive IMDA-aligned assessment with a gated red/amber/green posture, an AI risk appetite, and a board-ready report, plus an IMDA-aligned AI policy, an AI Verify self-assessment mapping, a PDPA-for-AI checklist and a board report template. kennedyadvisory.co

Kennedy Advisory provides independent, board-level AI governance advisory. This guide is a governance aid, not legal advice; Singapore's AI frameworks are voluntary and the PDPA is binding law — confirm your obligations with your advisers and regulators. Framework details current as at July 2026.