

The AI Ethics Review

A working toolkit for boards and design teams: the questions that catch what the outcomes test misses — and the frameworks taking legal effect.

Michael Kennedy · Former Group CIO, Robert Walters · Singapore / Global · July 2026

WHY THIS TOOLKIT

The outcomes test is not the whole question.

Most AI governance begins with a practical question: will the system work, and do the benefits outweigh the harms? That question matters — but even a careful weighing of outcomes can authorise systems a good organisation later regrets. There is a harder question underneath it: **who, in this design, is being treated merely as a means rather than also as an end?**

A design can be efficient, profitable and broadly beneficial, and still fail that test. This toolkit turns the principle into working review disciplines — usable in an architecture review, a vendor assessment or a board pack — and maps the regulatory frameworks now giving it legal force. It accompanies the essay series *The Duty of Care*.

The one-line test: the architectures you authorise should be ones you would be willing to live inside yourself.

DISCIPLINE 01

The asymmetry test

When a design works only because of an asymmetry — protected requesters and exposed targets, watchers who are themselves unwatched — treat that as an alarm. Asymmetry is where someone is being used without agency, transparency or recourse.

For every system that scores, ranks, filters or monitors people, ask:

- Knowledge** — does the person know the system is being used on them?
- Understanding** — can they understand what materially influenced the decision, in terms they could act on?
- Challenge** — is there a real route to contest an error, with someone empowered to fix it?
- Human authority** — is there a human with the authority, and the responsibility, to intervene and override?

The pattern to catch:

- Automated hiring is the canonical case: a candidate scored and filtered by a model they cannot inspect, on criteria they cannot challenge, for an efficiency whose benefits they never share. The imbalance — not the technology — is the thing to govern.
- Without the four protections above, efficiency is being purchased with someone else's agency.

DISCIPLINE 02

Five review disciplines

1. Use the categorical imperative as an architecture-review test

- Before shipping, ask: would a world in which every organisation deployed this feature still cohere — or would it quietly destroy the trust that makes the feature valuable?

2. Treat data as intent, not exhaust

- Logs, telemetry and query histories are personal data, not engineering by-products. A search can be an expression of fear, want or need. Value all three of useful, revealing and personal — not just the first.

3. Ask who carries the weight

- Average outcomes hide uneven harm. Identify which specific groups bear disproportionate cost — and which are quietly subsidising everyone else's benefit.

4. Make ethics evidential

- Trust accrues through demonstrable performance: documented decisions, audit trails — and the proposals you rejected, kept on the record.

5. Refuse the lowest local standard

- The strongest pull in any global programme is towards the weakest local requirement. Set one defensible common standard. Dignity is not jurisdictional.

THE LEGAL LAYER

The frameworks map — APAC first

Ethics leads; regulation follows. The law is the floor, not the standard — but the floor is rising. Status as at July 2026.

Framework	What it is / when it bites	What to do now
Singapore Model AI Governance Framework + AI Verify (IMDA/PDPC)	Pragmatic, board-friendly guidance and a testing toolkit; the reference point for responsible AI across much of APAC. Voluntary, but increasingly expected.	Use it as your internal review vocabulary; run AI Verify on one material use case as a proof of governance.
PDPA (Singapore) / GDPR Art. 22 (EU)	Data protection law already constrains automated decisions with significant effects; GDPR grants qualified protection from solely automated decisions.	Map which decisions are already in scope today — most organisations are surprised.
ISO/IEC 42001 — AI Management Systems	The certifiable management-system standard for AI — structurally the ISO 27001 of AI governance.	If you run 27001, you have the muscle memory. Gap-assess now; certify when a customer or regulator first asks.
EU AI Act	In force since 2024: prohibitions and AI literacy applied Feb 2025; general-purpose AI rules Aug 2025. The 2026 Digital Omnibus deferred high-risk (Annex III) obligations to 2 Dec 2027 and embedded-product AI to Aug 2028. Extraterritorial, GDPR-style.	Treat the deferral as a grace period, not a reprieve: classify your use cases against Annex III now, while it is cheap.
NIST AI RMF (US)	Voluntary risk-management framework; increasingly referenced in US contracts and procurement.	Adopt its vocabulary (map, measure, manage, govern) for anything US-facing.

The compliance trap: re-using customer data for an adjacent purpose may well be lawful; whether it respects the context in which the data was given is a separate question — and the honest answer is often no. Compliance keeps you out of trouble. Ethics earns you the right to operate data systems at scale.

IN THE REVIEW

The one-page ethics review

Ten questions to put to any AI system before it ships — or any vendor before you sign. Evidence, not assurance.

- Can we name the specific people this system scores, ranks, filters or watches?
- Do those people know, and can they understand what influenced the outcome?
- Is there a working route to challenge an error — tested, not just documented?
- Who is the named human with authority to override or switch it off?
- Would this design still work if every counterparty deployed the same thing at us?
- What personal intent is embedded in the data we are collecting as “telemetry”?
- Which group carries disproportionate cost, and did anyone ask them?
- Where is the written record of the designs we rejected, and why?
- Which jurisdiction’s standard are we actually applying — the best, or the cheapest?
- Would the executive sponsor be willing to live inside this system themselves?

ABOUT THE AUTHOR

Michael Kennedy is the founder of Kennedy Advisory and former Group CIO of Robert Walters, where he led a 150+ person technology organisation across 14 countries, delivered £3M+ in annualised savings, achieved ISO/IEC 27001 across all jurisdictions and took AI safely into production. He is CISSP certified and based in Singapore with a global remit.

Working through this for real?

Independent judgement on technology, cyber risk and AI — as an interim or fractional CIO, or at the board table. Engagements start with a conversation, not a proposal.

michael@kennedyadvisory.co · kennedyadvisory.co · Singapore / Global

© 2026 Kennedy Advisory, Singapore (UEN 53527056M). Companion to the essay series “The Duty of Care”. Regulatory positions as at July 2026 — verify before relying. You are welcome to share this document in full, with attribution. v1.0 — July 2026.