The Board Director's

AI Governance Guide

15 Practical Tools for AI Oversight Without Technical Expertise

July 2026

Impressum

IT AI Advisory Sdn. Bhd. (202601023072)

B-3A-8, Icon City, Jalan SS8/39, 47300 Petaling Jaya, Selangor, Malaysia

Responsible for content: Reto Gruenenfelder

reto@itaiadvisory.com · itaiadvisory.com

Important Notice

Purpose of this Guide

This guide is structured around questions, not directives. Its purpose is to create awareness and prompt board-level thinking — not to prescribe actions, guarantee outcomes, or replace independent professional judgment. No reader should treat any question or framework in this guide as a specific recommendation applicable to their organisation.

General Information Only

This guide is published information, not advice. It does not constitute advice of any kind — legal, financial, regulatory, compliance, or otherwise — and no advisory relationship of any nature is created by reading it. It is directed at professional readers, including board directors and senior executives of organisations that deploy or are considering deploying artificial intelligence. Any content authored by Reto Gruenenfelder is published in his capacity as founder and director of IT AI Advisory Sdn Bhd and not in any separate personal capacity.

Scope Limitation

This guide is not exhaustive. The frameworks, questions, and regulatory references it contains are illustrative starting points for board-level discussion. They do not constitute a complete or definitive statement of all AI governance obligations, regulatory requirements, or risk categories applicable to any organisation. The absence of a topic, framework, jurisdiction, or obligation from this guide does not imply it is inapplicable or immaterial. Boards and executives must assess their specific obligations independently.

No Regulated Financial Advice

For the avoidance of doubt, and without limiting the general principle above, this guide does not constitute regulated financial, investment, or compliance advice under any applicable financial services legislation — including but not limited to the laws of the member states of the Association of Southeast Asian Nations (ASEAN), the European Union, the United Kingdom, Switzerland, the United States of America, or any other jurisdiction. Nothing in this guide should be relied upon as a substitute for regulated advice from a licensed professional in the relevant field and jurisdiction.

No Warranty of Accuracy

This guide was published in July 2026. Information reflects the position as of that date only. No representation or warranty, express or implied, is made as to the accuracy, completeness, or timeliness of the information contained herein. Regulatory frameworks, laws, and standards referenced — including but not limited to the EU AI Act, MAS AI Governance Framework, SC Malaysia Technology Risk Management Guidelines, BNM Technology Risk Management Guidelines, FINMA Operational Risk and Resilience Guidelines, NIST AI Risk Management Framework, and ISO/IEC 42001 — are subject to ongoing development and change. Readers must verify the current status of any framework before relying on it.

Limitation of Liability

To the fullest extent permitted by applicable law, IT AI Advisory Sdn Bhd and its directors, officers, employees, and agents — including Reto Gruenenfelder acting in his capacity as founder and director — exclude all liability for any direct loss, indirect loss, consequential loss, loss of profit, loss of data, costs, expenses, or claims by third parties arising directly or indirectly from reliance on this guide. This exclusion does not apply to liability that cannot be excluded under mandatory applicable law, including liability for fraud or fraudulent misrepresentation. In particular, nothing in this disclaimer excludes or limits liability for damages arising from injury to life, body, or health, or for intentional or grossly negligent conduct, as required under the mandatory provisions of applicable law in any relevant jurisdiction.

No Professional Relationship

This guide does not create a professional relationship between IT AI Advisory Sdn Bhd, its directors (including Reto Gruenenfelder), or any associated persons, and the reader. Readers should seek independent professional advice appropriate to their specific circumstances and jurisdiction before making any governance, legal, regulatory, or business decisions.

Governing Law

This disclaimer and the guide, and any disputes arising from or in connection with them, are governed by and construed in accordance with the laws of Malaysia. The courts of Malaysia shall have non-exclusive jurisdiction to settle any such dispute, subject to any mandatory jurisdiction requirements of your own jurisdiction. Nothing in this disclaimer limits any rights you may have under the mandatory laws of your own jurisdiction.

Severability

If any provision of this disclaimer is found to be invalid, unlawful, or unenforceable under the laws of any jurisdiction, that provision shall be severed from the remainder of this disclaimer, which shall continue in full force and effect. The invalidity or unenforceability of any provision in one jurisdiction shall not affect its validity or enforceability in any other jurisdiction.

Publisher

IT AI Advisory Sdn. Bhd. is a commercial advisory firm. This guide is published free of charge as part of its professional outreach. The company offers paid advisory services to boards and executive teams; details are available at itaiadvisory.com. For information on how personal data is collected and processed, please refer to the Privacy and Cookie Policy available at itaiadvisory.com.

Views

The views expressed in this guide do not represent the position of any organisation with which Reto Gruenenfelder is or has been associated other than IT AI Advisory Sdn. Bhd.

Why This Guide Exists

Boards are increasingly expected by regulators to govern AI with the same rigour they apply to financial risk. Yet most directors have never received a practical framework for doing so.

This guide fills that gap. It is not a technical manual. It does not require you to understand how AI models work. It is a governance resource built around the questions boards need to ask, the risks they need to see, and the structures they need to put in place.

While the frameworks apply broadly to any organisation deploying AI, the guide is written primarily for financial services boards — where regulatory density, systemic risk, and data sensitivity make AI governance most consequential.

> *Boards don't need to understand how AI works. They need to understand who is accountable when it doesn't.*
>
> *— Reto Gruenenfelder*

The guide is structured in five parts, moving from foundations to action. Each section includes a board question you can take directly into your next governance discussion.

Contents

PART 1

AI Governance Foundations

Most boards are asking the wrong AI question. They ask: "What AI should we invest in?" The right question is: "Do we have the governance structures to oversee AI we've already deployed?" In most organisations, the answer is no.

AI pilots that launched in 2022–2024 are becoming production systems without proper governance, monitoring, or board visibility. The governance collapse will follow.

Framework 1 — Before Approving Any AI Initiative

Five questions boards are encouraged to ask before approving any AI project:

1. **Accountability**
   Who is accountable if this AI system causes harm—to customers, staff, or the market?
2. **Explainability**
   Can management explain how this AI makes its decisions — in plain language?
3. **Audit process**
   What is our process for auditing AI performance over time, not just at launch?
4. **Regulatory exposure**
   Have we mapped our regulatory exposure—including applicable AI governance regulations and sector-specific frameworks—in all jurisdictions where we operate?
5. **Risk framework**
   Does our current risk framework cover AI-specific risks, or are we applying old models to a new problem?

Most boards get presented the opportunity. Very few get presented the governance gap alongside it.

> **Board Question**
> *Can management answer all five questions for every AI initiative currently awaiting board approval?*

Framework 2 — Why AI Transformation Fails

Every large-scale transformation that struggled had the same root cause: not technical failure, but governance failure. The same pattern is emerging in AI:

1. **No clear ownership**
   AI initiatives sit between IT, data science, and the business. Nobody is fully accountable.
2. **Pilot without purpose**
   Teams launch AI experiments without defined success metrics or board-level business cases.
3. **Risk misclassified**
   AI risk gets treated as a technology risk when it is actually operational, reputational, and regulatory.
4. **Board kept at arm's length**
   Management presents AI as too technical for board oversight. Boards accept this framing.
5. **No feedback loop**
   AI systems are deployed and forgotten. No monitoring. No audit. No course correction.

The fix isn't a better AI model. It's a better governance structure.

> *The most dangerous moment in any transformation wasn't the launch. It came later — when the original sponsors had moved on, the budget was under pressure, and the team was left without clear executive ownership.*

Board Question
Which AI systems in your organisation are now in production — and who owns them?

PART 2

New Risk Categories: Generative AI and Explainability

Framework 3 — Generative AI: A Different Category of Board Risk

Traditional AI makes decisions. Generative AI generates content. That difference sounds technical. For boards, it is a different category of risk entirely.

A biased credit model affects the people it scores. A GenAI system that hallucinates, misrepresents, or produces harmful content affects everyone who reads it — customers, regulators, markets, the press. Brand exposure, legal liability, and reputational damage can arrive in hours, not quarterly risk reports.

Boards that approved "AI strategy" in 2024 are not automatically covering the GenAI systems their teams deployed in 2025. Generative AI is a content engine operating at scale, in your organisation's name.

The five questions below complement Framework 1's pre-approval checklist—they are not a replacement for it. Apply both to any AI initiative that includes a generative component.

Five questions boards are encouraged to ask before any GenAI system goes live:

1. **Output review**
   Who reviewed the outputs before deployment — and on what criteria?
2. **Human in the loop**
   What happens when the system produces something factually wrong? Is there a human in the loop, and at what point?
3. **Data rights**
   What data was used to train or fine-tune this model — and do we have the rights to it?
4. **Scope testing**
   How does this system behave when asked questions outside its intended scope? Has that been tested?
5. **Production ownership**
   Who owns this system in production — not just at launch? Who gets called at 7 am?

Board Question
Do you know which generative AI systems in your organisation are currently customer-facing?

Framework 4 — The XAI Test: Explainability as Accountability

A board cannot govern what it cannot interrogate. A regulator cannot audit what cannot be explained. Explainability is not a feature you add later — it is a governance requirement you build in from the start.

Three questions that determine whether your AI system meets the explainability standard:

1. **The customer test**
   Can you explain this decision in plain language to the person it affected? Could a customer understand why the credit was declined, the claim was rejected, or the application was flagged?
2. **The regulator test**
   Can you explain this specific decision — on this specific input, producing this specific output — to your regulator? Is there an audit trail?
3. **The board test**
   Can management articulate clearly and simply how the system works and what its known limitations are — without technical slides?

If the answer to any of these is no, the system has an explainability gap. That gap is a governance gap. And applicable governance frameworks and regulatory guidance hold that the board takes responsibility for closing it.

> **Board Question**
> *Which AI systems in your organisation could not explain their decisions to your regulator today?*

PART 3

The Regulatory Landscape in 2026

AI regulation is moving faster than most boards realise. Seven major frameworks are now operational simultaneously. Boards that have not formally reviewed their exposure are already behind.

Framework 5 — The AI Regulatory Map

Framework	Scope	Key Board Obligation
EU AI Act	Europe + any organisation serving EU customers	Risk-based classification. High-risk AI systems require conformity assessments, technical documentation, and human oversight mechanisms.
FCA/PRA AI Guidance	United Kingdom — FCA and PRA regulated firms	Principles-based. FCA AI Strategy and PRA model risk supervisory statement (SS1/23) set governance, explainability, and accountability expectations for AI-driven decisions. UK firms with EU operations must also assess their exposure to the EU AI Act separately.
MAS AI Governance Framework	Singapore (widely referenced across ASEAN)	Principles-based. Covers internal governance, risk management, and consumer protection. Referenced by regulators across ASEAN.
SC Malaysia Technology Risk Management Guidelines	Malaysia's capital markets	Algorithm governance, model risk, and technology controls for capital market intermediaries.
BNM Technology Risk Management Guidelines	Malaysia — banks, insurers, payment system operators	Technology risk governance for financial institutions regulated by Bank Negara Malaysia. Covers AI and algorithmic decision-making within the broader technology risk framework. Applies to licensed banks, insurance companies, and payment service providers.
ISO/IEC 42001	Global standard	The global AI management system standard. Most management teams are aware of it. Few boards have formally reviewed it or received a briefing on what it requires of the organisation.
FINMA Operational Risk & Resilience Guidelines	Switzerland — FINMA-regulated banks, insurers, and financial market infrastructure	Operational and technology risk governance for FINMA-regulated institutions. AI systems fall within the technology risk framework under Circular 2023/1. Swiss institutions with EU operations must also assess their exposure to the EU AI Act.
NIST AI RMF	US + global reference	Risk management framework for AI systems published by the US National Institute of Standards and Technology. Widely referenced by regulators and auditors outside the US as a technical baseline for AI risk governance.

Board Question

Which of these regulatory frameworks has your board formally reviewed in the last 12 months?

Framework 6 — Evaluating an AI Vendor Pitch

Before approving any significant AI investment, ask management to answer these six questions. If they cannot answer all six, the investment is not ready for board approval.

1. **Data rights**
   What data does this system require—and do we have the right to use it?
2. **Failure response**
   What happens when the model is wrong? Who is notified, and how fast?
3. **Independent audit**
   Has an independent third party audited this vendor for bias and accuracy?
4. **Exit strategy**
   What is our exit strategy if we need to switch vendors at short notice?
5. **Regulatory class**
   Does this system fall under a high-risk classification under any AI or financial services regulation applicable to our operating jurisdictions?
6. **Production ownership**
   Who in our organisation will own this system in production — not just at launch?

> **Board Question**
> *Can management answer all six questions before your next AI investment decision?*

PART 4

The Board's Role and Committee Structure

Boards don't need an AI expert director. They need directors who know which questions management cannot answer—and are willing to ask them out loud. Technical expertise is nice to have. Governance courage is essential.

> *"Understanding is not the prerequisite for governance. Accountability is."*
>
> *— Reto Gruenenfelder*

Framework 7 — Five Governance Myths

1. Myth: "AI governance is an IT responsibility."
Reality: IT manages AI infrastructure. The board governs AI risk, ethics, and accountability.

2. Myth: "We need a technical expert on the board."
Reality: Boards need governance competence, not technical expertise. The questions that matter are not technical.

3. Myth: "Our AI projects are too small to need governance."
Reality: Small AI projects become large ones. In practice, retrofitted governance on a running system is far more costly than governance built in from the start.

4. Myth: "Data protection compliance covers our AI obligations."
Reality: Data protection law — including PDPA in Malaysia (as amended in 2024), PDPC requirements in Singapore, GDPR in Europe, and equivalent frameworks across ASEAN and beyond — covers personal data. AI governance covers much more: model accountability, systemic risk, market conduct, and the full scope of applicable AI-specific regulation.

5. Myth: "We have a Chief AI Officer — the board is covered."
Reality: A CAIO manages AI execution. The board governs AI risk. These are not the same job.

> **Board Question**
> *Which of these five myths is most embedded in how your board currently approaches AI?*

Framework 8 — Three Roles Every AI Governance Committee Needs

1. The Challenger
Reads every management AI briefing with scepticism. Asks what's been left out. Requests the risk register, not just the roadmap.

2. The Translator
Moves between technical language and governance language. Doesn't need to build AI — needs to explain it to the full board in plain terms.

3. The Regulator-Watcher
Tracks applicable AI regulation, governance standards, and sector-specific guidance across all

jurisdictions the organisation operates in. Ensures the board is briefed before regulatory changes create exposure, not after.

You don't need three separate people. But you need all three roles covered.

> **Board Question**
> *Are all three roles — Challenger, Translator, and Regulator-Watcher — covered in your current AI governance structure?*

Framework 9 — Workforce Transformation and AI Literacy

Governing AI means governing what AI does to the people in your organisation — and what it requires of them. Three questions sit at the heart of this: where human judgment should remain, what we are doing about the workforce we are transforming, and how we build organisation-wide AI literacy.

1. **Human-in-the-loop decisions**
   High-impact decisions — credit approvals, insurance assessments, employment screening, access to services — require clear protocols defining where AI advises and where a human decides and is accountable. Applicable governance frameworks and regulatory guidance suggest boards should require management to document this, review it regularly, and verify alignment with applicable regulatory requirements.

2. **Workforce transformation**
   AI adoption will change roles, headcounts, and skill requirements across the organisation. Boards are encouraged to require management to present a workforce transformation strategy: not just which roles are affected, but how the organisation is investing in reskilling, redeployment, and new role creation. Displacement without a plan is a material governance gap.

3. **AI upskilling for the whole organisation**
   A workforce that does not understand the AI systems it works alongside cannot flag when those systems go wrong. Applicable governance frameworks and regulatory guidance suggest boards expect management to run systematic AI literacy programmes across the organisation — not just for technical staff. Customer-facing employees, compliance teams, risk functions, and senior management all need sufficient understanding to work with AI responsibly. Upskilling is not an HR initiative. It is a governance requirement.

> **Board Question**
> *Has management mapped which decisions in your organisation require human accountability — and committed to a concrete AI upskilling plan for the workforce?*

Framework 10 — Five AI Risks Regulators Are Watching

1. Algorithmic discrimination
AI systems that produce systematically different outcomes for groups defined by protected characteristics under applicable law — including but not limited to race, ethnicity, gender, age, religion, disability, and other immutable or legally recognised attributes.

2. Model opacity
Systems whose decisions cannot be explained to customers, regulators, or courts. This is increasingly unacceptable under AI regulation globally and may constitute a compliance failure in jurisdictions with explainability requirements.

3. Data quality failure
AI trained on incomplete, biased, or outdated data that produces unreliable outputs at scale.

4. Third-party dependency
Organisations using AI systems they don't control and don't fully understand — including but not limited to those embedded in vendor software, cloud platforms, API services, or other third-party integrations.

5. Governance vacuum
AI systems in production with no named owner, no monitoring cadence, and no escalation protocol. The most common failure mode of all.

> **Board Question**
> *Which of these five risks does your board have full visibility on?*

PART 5

Building Your AI Governance Framework

A board AI governance framework is not a policy document that sits in a drawer. It is a living structure. Here is what it actually looks like in practice.

Framework 11 — The Four-Component Board Framework

The following represent minimum governance components. Organisations should assess what additional requirements apply based on their specific risk profile, regulatory environment, and AI maturity.

1. **Oversight mandate**
   A board resolution formally assigning AI governance responsibility — to the full board, a committee, or a named director. On record. Reviewable.
2. **Reporting cadence**
   Management reports to the board on AI systems in production at a minimum quarterly frequency. Includes: new deployments, incidents, regulatory developments, and model performance.
3. **Risk appetite statement**
   A written board-approved statement of which AI risks the organisation will accept, mitigate, or refuse. Not written by management alone.
4. **Escalation protocol**
   A defined path for AI-related incidents to reach the board promptly. Not discovered in a press report. The specific escalation timeframe should reflect the organisation's regulatory obligations and risk profile.

> **Board Question**
> *Does your board have all four components of a real AI governance framework in place?*

Framework 12 — How to Run Your First Board AI Audit

This doesn't require external consultants or technical expertise. It requires a board with the will to ask. Most boards that complete Step 1 are surprised by what they find.

1. **Step 1: Inventory**
   Ask management to produce a complete list of AI systems currently in production. Every one. Including those managed by vendors or embedded in third-party software.
2. **Step 2: Risk classification**
   For each system: what decisions does it influence? What happens if it's wrong? Which regulatory frameworks apply?
3. **Step 3: Governance check**
   For each high-risk system: who owns it? When was it last reviewed? What monitoring is in place? Has it ever produced an unexpected outcome?

4. **Step 4: Gap report**
   Where are the accountability gaps? What has no owner? What is not monitored? What has not been reviewed?

5. **Step 5: Board resolution**
   Take the gap report to the full board. Assign a named owner to each identified gap. Set a remediation deadline and a date for the board to review progress. A gap report that does not produce a resolution is not governance — it is documentation. Where the board disputes a risk classification or disagrees on prioritisation, that disagreement should itself be recorded — with a date by which it will be reviewed.

Board Question

Has your board completed Step 1 — a full inventory of every AI system currently in production, including those embedded in vendor software?

Framework 13 — Real vs. Performance Governance

Ask: "How would we know if our AI governance was not working?" The answer tells you everything.

Performance Governance	Real Governance
Has a policy document	Has a named owner
Approved once, never revisited	Reviewed on a defined cadence
Risk log not updated	Risk log reviewed and updated regularly
Board briefed annually	Board briefed quarterly
Monitoring delegated to IT with no board reporting	AI incidents escalated to the board within a defined timeframe

Board Question

Honestly assessed: which column describes your organisation's AI governance today?

Framework 14 — AI Scenario Planning

Leading boards apply scenario exercises to AI — the same way they apply them to financial stress events and cyber incidents. Three scenarios to run:

1. **Discrimination scenario**
   "What happens if our most important AI system produces a discriminatory outcome — and it's in tomorrow's newspaper?"
2. **Regulatory freeze**
   "What do we do if a regulator freezes our AI-based credit decisions pending investigation?"
3. **Industry contagion**
   "What is our response if a competitor's AI failure triggers industry-wide regulatory scrutiny?"

These scenarios are not hypothetical. They are happening.

This list is not exhaustive. Boards should also consider: an adversarial attack or model poisoning event that corrupts outputs without triggering obvious alerts; AI-enabled fraud perpetrated against the organisation using its own customer data; and a third-party AI vendor failure that disrupts a critical business process with no tested fallback.

> **Board Question**
> *Has your board run any of these scenarios — and if not, when will you schedule the first one?*

Framework 15 — Financial Services: Five Differentiating Factors

The following represent key differentiating factors in financial services AI governance. They are not exhaustive — organisations should assess additional factors specific to their business lines, regulatory obligations, and operating jurisdictions.

1. **Regulatory density**
   Banks, insurers, and exchanges operate under layered oversight. AI governance should integrate into the existing compliance architecture—not sit beside it.

2. **Systemic risk exposure**
   A biased credit model or flawed trading algorithm doesn't just harm one customer. It can cascade.

3. **Data sensitivity**
   Financial AI operates on some of the most sensitive personal and commercial data. Every AI deployment is simultaneously a data governance question.

4. **Speed vs control**
   Fintechs operate under the same regulatory obligations as established institutions — but without legacy systems, they can build AI governance in from the start. Incumbents will need to retrofit governance onto existing infrastructure, processes, and cultures not designed with AI in mind. This creates a structural disadvantage that the board should actively address, not use as an excuse for delay.

5. **Trust as the product**
   In financial services, trust is the product. One AI governance failure can erase years of brand equity.

Board Question
Does your board treat AI governance as structurally integrated into your compliance architecture, or as a separate initiative alongside it?

What Comes Next

The boards that govern AI well in 2026 are not the most technically sophisticated. They are the most intellectually honest — willing to slow down a high-profile AI initiative when the governance is not ready, even under commercial pressure.

Boards that govern AI well don't have more knowledge than others. They have better conversations. It just requires someone at the board table willing to start.