01 · DIAGNOSTIC · 10 QUESTIONS

# Make AI governable.

A 10-question test every Australian board should run before its next risk committee.

*Three or more REDs is a board-level conversation, not an IT one.*

MAPPED TO · AICD DIRECTOR-DUTY GUIDANCE · APRA CPS 230 · CPS 234 · AS ISO/IEC 42001

## Q01

**Do we know where AI is being used in our organisation today — including AI embedded inside third-party products and SaaS?**

WHY IT MATTERS Shadow AI is the most common governance blind spot. Most boards cannot answer this.

## Q02

**Has one named person or body been made accountable to the board for AI risk, oversight and incident response?**

WHY IT MATTERS Without a single accountable owner, AI risk defaults to technical teams and the board has no escalation path.

## Q03

**Does our risk appetite statement explicitly address AI risks — or is AI inheriting controls designed for traditional IT?**

WHY IT MATTERS CPS 230 and ISO 42001 expect AI-specific risk management. Traditional controls do not cover drift, hallucination or autonomous decisioning.

## Q04

**For our highest-risk AI use cases, who has authority to stop the system on a Friday afternoon — and when was the rollback last tested?**

WHY IT MATTERS Every production AI system needs a named owner, a documented kill-switch and a tested rollback. Most do not.

## Q05

**What is our policy on directors and the executive team using public generative AI tools for organisational work?**

WHY IT MATTERS AICD has explicitly red-flagged public AI used on board papers. Most Australian organisations have no written policy.

## Q06

**How does our procurement framework assess AI risk embedded in vendor products, SaaS contracts and supply chain?**

WHY IT MATTERS Most AI risk now enters through the supply chain — not in-house builds. Procurement was not designed to evaluate vendor AI.

## Q07

**What evidence do we have that the data feeding our AI systems is accurate, lawful, current and free of material bias?**

WHY IT MATTERS AI inherits the quality of its inputs. Without documented lineage, every output is unprovable.

## Q08

**What AI training has the board itself received — and how recently?**

WHY IT MATTERS AICD has tied director duties directly to AI literacy. A board that cannot interrogate an AI risk paper is an exposed board.

## Q09

**How is the use of AI to draft, summarise or analyse board papers, minutes and committee documents currently governed?**

WHY IT MATTERS AICD's most recent publication flags AI-generated minutes as discoverable. Summaries of board packs risk legal privilege.

## Q10

**How would we explain — to a regulator or to a customer — a decision our AI system made yesterday?**

WHY IT MATTERS The transparency test. Australian AI Ethics Principles, ISO 42001 and regulator expectations all assume you can answer this.

02 · SCORING RUBRIC

# Score yourself honestly.

A 30-minute working session with your chair, company secretary and one risk-committee member. *Score each question RED, AMBER or GREEN using the rubric below.*

| # | RED | AMBER | GREEN |
|---|---|---|---|
| Q01 | No AI inventory exists. | Inventory exists for internally-built AI only. | Live register covers built, procured, and embedded AI. |
| Q02 | Accountability sits in IT or is undefined. | Named owner without formal charter or reporting line. | Charter, RACI and reporting cadence agreed at board level. |
| Q03 | AI risk is not in the risk appetite statement. | AI captured under generic technology or cyber risk. | AI-specific appetite, controls and escalation thresholds. |
| Q04 | No identified owner, no rollback plan. | Owner named, rollback documented but untested. | Both in place and tested in the last six months. |
| Q05 | No policy, or informal "we don't allow it". | Generic acceptable use policy, not AI-specific. | AI-specific protocol with vetted tools and training. |
| Q06 | No AI-specific clauses or vendor due diligence. | AI risk assessed for new vendors only. | AI risk built into onboarding and annual re-certification. |
| Q07 | No data lineage tracking for AI inputs. | Lineage tracked for some systems but not consolidated. | Documented lineage and bias testing for material use cases. |
| Q08 | None. | One-off briefing in the last 24 months. | Annual board AI training refreshed on regulatory change. |
| Q09 | No policy. Ad-hoc use. | Topic discussed at board, no documented protocol. | Written protocol on AI use of board material. |
| Q10 | We could not. | Partial logs but not consolidated or customer-facing. | Audit trail, explainability framework and remediation path. |

**3+ REDs**
Material exposure under CPS 230 and AS ISO/IEC 42001. Time for a structured assessment.

**4+ AMBERs**
Foundational work in place but un-validated. Pressure-test before regulators do.

**Mostly GREEN**
Ahead of the AU mid-market. Focus on board-level AI use and director literacy.

MYCE|IUM™

AI governance, security and agents for Australian boards and regulated industries. The 4-week AI Security Health Check™ scores your organisation across 14 dimensions and 4 pillars — mapped to APRA CPS 230, CPS 234 and AS ISO/IEC 42001.

THE FIRST CALL IS DIAGNOSTIC

*Sixty minutes. No pitch. No proposal until it makes sense.*

Sara Shakib · **admin@themyceliumgroup.com.au**

+61 401 844 836 · themyceliumgroup.com.au

SOURCES *AICD & HTI, A Director's Guide to AI Governance (2024); AI Governance Checklist for SME and NFP Directors (2024); Effective Board Minutes and the Use of AI (2025); AI Use by Directors and Boards: Early Insights (2025).*