01 · DIAGNOSTIC · 10 QUESTIONS · BANKING & MUTUALS

Make AI governable.

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

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

MAPPED TO · AICD · APRA CPS 230 · CPS 234 · CPS 220 · AFSL · AS ISO/IEC 42001

Q01

Do we know where AI is being used in our organisation today — including AI embedded inside our core banking system, third-party products, payments, fraud and KYC platforms?

WHY IT MATTERS Shadow AI is the most common governance blind spot. In banking, embedded AI inside the core, vendor APIs and SaaS tools is rarely visible to the board.

Q02

Has one named person or body been made accountable to the board for AI risk, oversight and incident response — and is this reflected in our CPS 230 operational risk profile?

WHY IT MATTERS CPS 230 expects clear accountability for operational risk. 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 under CPS 220?

WHY IT MATTERS CPS 230, CPS 220 and ISO 42001 expect AI-specific risk management. Traditional IT controls do not cover model drift, hallucination or AI-driven customer decisioning.

Q04

For our highest-risk AI use cases (credit decisioning, fraud, AML, customer servicing), who has authority to stop the system — 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. For customer-facing decisioning, kill-switch authority must be operable inside business hours.

Q05

What is our policy on directors, executives and customer-facing staff using public generative AI tools — including drafting customer communications, advice or board papers?

WHY IT MATTERS AICD has explicitly red-flagged public AI used on board papers or confidential material. APRA-regulated firms have heightened sensitivity around customer and credit data.

Q06

How does our procurement framework assess AI risk embedded in core banking, payments, fraud-detection, KYC and SaaS vendor products?

WHY IT MATTERS Most AI risk now enters through the supply chain. CPS 230 requires material service provider oversight — AI embedded inside vendor systems sits squarely inside that obligation.

Q07

What evidence do we have that the data feeding our AI systems (especially credit, fraud and customer-decisioning models) is accurate, lawful, current and free of material bias?

WHY IT MATTERS AI inherits the quality of its inputs. Without documented lineage and bias testing, every customer outcome is unprovable to ASIC, AFCA or a class-action solicitor.

Q08

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

WHY IT MATTERS AICD has tied director duties directly to AI literacy. For APRA-regulated boards, this expectation is elevated.

Q09

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

WHY IT MATTERS AICD's most recent publication flags AI-generated minutes as discoverable. For regulated entities, AI summaries of board packs risk legal privilege and regulator scrutiny.

Q10

How would we explain — to APRA, ASIC, AFCA or a customer — a decision our AI system made yesterday? (Especially a credit, fraud or AML decision.)

WHY IT MATTERS The transparency test. APRA, ASIC and AFCA all expect that automated decisioning can be explained. Customer remediation depends on it.

02 · SCORING RUBRIC · BANKING & MUTUALS

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 core banking, vendor and embedded AI — refreshed quarterly.
Q02	Accountability sits in IT or is undefined.	Named owner without formal charter or CPS 230 reporting.	Named owner, charter, RACI and CPS 230-aligned reporting 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, integrated with CPS 220.
Q04	No identified owner, no rollback plan.	Owner named, rollback documented but untested.	Both in place and tested in the last six months. Customer impact path documented.
Q05	No policy, or informal “we don’t allow it”.	Generic acceptable use policy, not AI-specific.	AI-specific protocol with vetted tools, training, and customer-data guardrails.
Q06	No AI clauses for core banking or material suppliers.	AI risk assessed for new vendors only.	AI risk built into onboarding, annual re-certification and CPS 230 material service provider oversight.
Q07	No lineage or bias testing for credit/fraud models.	Lineage tracked for some systems but not consolidated.	Documented lineage and bias testing for material decisioning models.
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, including privilege and retention.
Q10	We could not.	Partial logs but not consolidated or customer-facing.	Audit trail, explainability framework and customer remediation path documented and tested.

3+ REDs

Material exposure under APRA CPS 230, CPS 234 and AS ISO/IEC 42001. Prudentially significant. Time for a structured assessment.

4+ AMBERs

Foundational work in place but un-validated. Pressure-test before your next APRA prudential review.

Mostly GREEN

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

MYCE | IUM™

AI Security Health Check™ — BANKING & MUTUALS Edition. The 4-week diagnostic scores your organisation across 14 dimensions and 4 pillars — sector-tuned and mapped to the standards your regulator reads.

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); APRA CPS 230 (Operational Risk Management), CPS 234 (Information Security), CPS 220 (Risk Management); AICD, Effective Board Minutes and the Use of AI (2025).*