

Make AI governable.

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

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

MAPPED TO • AICD • SOCI ACT • CIRMP • AEMO • AESCSF • AS ISO/IEC 42001

Q01

Do we know where AI is being used in our organisation today — including AI embedded inside operational technology (OT), SCADA, EMS, metering and third-party SaaS?

WHY IT MATTERS Shadow AI is the most common governance blind spot. In energy, embedded AI inside OT, EMS and metering systems 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 Critical Infrastructure Risk Management Program?

WHY IT MATTERS The CIRMP requires named accountability for material risks. 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 and OT?

WHY IT MATTERS SOCI, AESCSF and ISO 42001 expect AI-specific risk management. Traditional IT/OT controls do not cover model drift, hallucination or AI-driven operational decisioning.

Q04

For our highest-risk AI use cases (grid optimisation, demand forecasting, asset health, market bidding), who has authority to stop the system during a grid event — 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 grid-impacting AI, kill-switch authority must hold under peak demand and stress conditions.

Q05

What is our policy on directors, executives and field staff using public generative AI tools — including operational reports, incident summaries and board papers?

WHY IT MATTERS AICD has explicitly red-flagged public AI used on board papers. For SOCI-covered entities, leakage of operational data to public AI is a notifiable concern.

Q06

How does our procurement framework assess AI risk embedded in OT vendor products, EMS contracts, metering systems and supply chain?

WHY IT MATTERS Most AI risk now enters through the supply chain — particularly through control-system vendors, demand-management platforms and metering. SOCI explicitly covers material supplier risk.

Q07

What evidence do we have that the data feeding our AI systems (especially demand forecasting, asset health and market bidding) is accurate, lawful, current and free of material bias?

WHY IT MATTERS AI inherits the quality of its inputs. Without documented lineage and operational validation, every operational decision is unprovable to AEMO or the AER.

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 SOCI-covered boards, this expectation is elevated.

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. Operational sensitivity adds an extra concern around AI processing of board material in energy boards.

Q10

How would we explain — to AEMO, the AER, the regulator or a customer — a decision our AI system made yesterday? (Especially one affecting grid operation, market bids or customer billing.)

WHY IT MATTERS The transparency test. AEMO, the AER, AESCSF and emerging regulator expectations all assume operational AI decisions can be explained.

02 · SCORING RUBRIC · ENERGY & UTILITIES

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 IT-side AI only.	Live register covers IT, OT, EMS, metering and embedded AI — refreshed quarterly.
Q02	Accountability sits in IT or is undefined.	Named owner without formal charter or CIRMP linkage.	Named owner, charter, RACI and CIRMP-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; AESCSF-aligned.
Q04	No identified owner, no rollback plan.	Owner named, rollback documented but untested.	Both in place and tested in the last six months. Grid-event procedure 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 operational-data guardrails.
Q06	No AI clauses for OT or material suppliers.	AI risk assessed for new vendors only.	AI risk built into onboarding, annual re-certification and SOCI material supplier oversight.
Q07	No lineage tracking for AI inputs.	Lineage tracked for some systems but not consolidated.	Documented lineage and operational validation 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 operational sensitivity.
Q10	We could not.	Partial logs but not consolidated or customer-facing.	Audit trail, explainability framework and AEMO/regulator-facing remediation path.

3+ REDS	4+ AMBERS	Mostly GREEN
Material exposure under the SOCI Act, AESCSF and AS ISO/IEC 42001. Critical-infrastructure significant. Time for a structured assessment.	Foundational work in place but un-validated. Pressure-test before your next CIRMP review.	Ahead of the AU utility mid-market. Focus on board-level AI use and director literacy.