ISO 42001 Readiness Checklist

A board-level self-assessment for Australian organisations preparing for an AI management system

NETEVO | netevo.com.au

v1.0 — July 2026

How to use this checklist

This is a self-assessment instrument for boards, executives, and risk leads who want to know how far their organisation is from being able to stand up an AI management system of the kind ISO/IEC 42001 describes. Work through the twenty-one questions below with the people who actually operate your AI systems — not in the abstract. Every item is a yes/no question, and every honest "no" marks a gap worth closing before an ISO/IEC 42001 program begins; closing gaps in the order the phases suggest is almost always cheaper than discovering them mid-program. This checklist is NETEVO's own operational instrument, informed by the intent of the standard but written entirely in our own words. It does not reproduce ISO/IEC 42001 and is not a substitute for it — organisations pursuing conformity should purchase the standard itself from ISO (www.iso.org) or, in its Australian adoption AS ISO/IEC 42001:2023, from Standards Australia (store.standards.org.au). This checklist explains readiness in operational terms; it is not legal advice and not a conformity assessment — consult your own advisers on how these instruments apply to your specific circumstances.

Phase 1 — Accountability & visibility

You cannot govern what nobody owns and nobody can see. Phase 1 asks whether the basic facts of your AI use are known, owned, and reported.

1. Is there a named individual — not a committee — accountable for AI systems across the organisation? Yes □ No □ *Why it matters: shared accountability is deniable accountability; a management system needs one name the board can ask.*

2. Do you hold a current inventory of every AI system in use, including AI embedded inside vendor products? Yes □ No □ *Why it matters: most organisations under-count — the AI inside your CRM, hiring tools, and productivity suite is still your exposure.*

3. Does that inventory capture unofficial or staff-initiated AI use, not just the systems IT procured? Yes □ No □ *Why it matters: ungoverned tools handling company data are usually the largest unpriced risk on the register.*

4. Is every inventoried system assigned a risk tier based on what it touches — customers, personal information, money, safety? Yes □ No □ *Why it matters: without tiering, a meeting-notes summariser and a credit-decisioning model receive the same scrutiny, which means the wrong one gets it.*

5. Does AI appear on a standing board or risk-committee agenda with a defined reporting line? Yes □ No □ *Why it matters: an escalation path invented during an incident is not an escalation path.*

6. Is there a written policy on acceptable AI use that staff can actually find, read, and apply? Yes □ No □ *Why it matters: a policy nobody can locate governs nobody; auditors ask staff, not the intranet.*

7. Do the people who operate or approve AI systems know they hold that role, and what it requires of them? Yes □ No □ *Why it matters: responsibility that exists only in an org chart fails at exactly the moment it is needed.*

Phase 2 — Control & gate

Phase 1 tells you what you have. Phase 2 asks whether anything stands between an AI system and the moment it starts acting on your customers, your data, and your name.

8. Is there an approval gate that every AI system must pass before it goes live — with authority to say no? Yes □ No □ *Why it matters: a gate that has never refused anything is a rubber stamp; the test is whether "not yet" is a possible outcome.*

9. For higher-tier systems, is an impact assessment completed before deployment — who could this harm, and how would you know? Yes □ No □ *Why it matters: assessing impact after go-live converts a design decision into an incident response.*

10. Are third-party and vendor AI arrangements reviewed for how your data is used, how the model changes over time, and who carries responsibility when it fails? Yes □ No □ *Why it matters: outsourcing the system does not outsource the accountability — your regulator and your customers still come to you.*

11. Can you state, for each significant system, what data feeds it and whether you hold the rights to use that data this way? Yes □ No □ *Why it matters: data provenance is where AI risk, privacy risk, and intellectual-property risk converge — and it is the question outside scrutiny reaches first.*

12. Are there defined points where a human reviews, intervenes in, or overrides the system's output — chosen by risk, not by convenience? Yes □ No □ *Why it matters: oversight placed where it is easy rather than where it is needed produces the appearance of control without the substance.*

13. Do staff know how to report an AI failure or suspected harm, and is there a defined path for who responds? Yes □ No □ *Why it matters: the gap between "someone noticed" and "someone acted" is where small failures become reportable ones.*

14. Do material changes — new model version, retrained data, altered prompts or configuration — go back through review before release? Yes □ No □ *Why it matters: an AI system approved once and modified freely is, after a few months, an unapproved system.*

Phase 3 — Operate & evidence

Controls that ran yesterday prove nothing about today. Phase 3 asks whether your governance operates continuously — and whether it leaves evidence an outsider could inspect.

15. Are live AI systems monitored for behaviour and performance against what was approved — not just for uptime? Yes □ No □ *Why it matters: infrastructure monitoring tells you the system is running; it does not tell you the system is still doing the right thing.*

16. Is there a scheduled review that asks whether each system's behaviour has drifted from its approved purpose? Yes □ No □ *Why it matters: models, data, and usage all shift quietly; drift found on a schedule is a finding, drift found by a customer is an incident.*

17. Could you show an independent auditor the record — approvals, assessments, incidents, reviews — without assembling it from inboxes first? Yes □ No □ *Why it matters: in any audit, evidence you cannot produce is treated as a control that does not exist.*

18. Are your AI controls themselves tested on an internal audit rhythm — someone checking that the gate, the monitoring, and the incident path actually work? Yes □ No □ *Why it matters: a management system is defined by checking itself; untested controls decay silently.*

19. When a review or incident produces a finding, is it tracked to closure — with someone accountable for closing it? Yes □ No □ *Why it matters: findings without a corrective loop are a growing list of things you knew about and did not fix — the worst possible position.*

20. Is there a defined path for retiring an AI system — including what happens to its data, its outputs, and its dependencies? Yes □ No □ *Why it matters: decommissioning is the control everyone forgets; abandoned systems keep running, and keep deciding, unowned.*

21. Does senior leadership periodically review the whole posture — inventory, incidents, findings, resourcing — and adjust it? Yes □ No □ *Why it matters: a management system that never changes in response to what it learns is paperwork, not management.*

Reading your result

Mostly yes (16 or more). The bones of an AI management system already exist in your organisation. The work ahead is consolidation: formalising what runs informally, connecting the pieces into one system, and hardening the evidence trail so that what you do is also what you can show.

Mixed (roughly 8 to 15). This is the most common position. Typically Phase 1 is largely in place, Phase 2 is partial, and Phase 3 is thin — governance exists on paper and in pockets, but does not yet operate as a system or leave inspectable evidence. Sequence the gaps in phase order: visibility first, gates second, operating evidence third.

Mostly no (fewer than 8). Start with Phase 1 and nothing else. A named owner, a real inventory, and a risk tiering are the prerequisites for every other control on this list — attempting gates and monitoring before you know what you are governing produces effort without coverage.

Whatever the band, the same principle applies: each “no” is a specific, closable gap, and the order above is the economical order to close them in.

What NETEVO does

NETEVO encodes governance obligations into executable controls with an evidence trail — the operating layer beneath a standard like ISO/IEC 42001, where policy becomes something a system enforces and an auditor can inspect. NETEVO does not certify anyone: certification is performed by accredited certification bodies.

Read next:

- What ISO 42001 asks of an Australian organisation, in operational terms — **netevo.com.au/insights/iso-42001-explained/**
- The full ANZ regulatory evidence base — **netevo.com.au/whitepapers/ai-governance-anz/**
- How NETEVO operationalises AI governance — **netevo.com.au/solutions/ai-governance/**
- Talk to us about your result — **netevo.com.au/contact/**

Gregory McKenzie Registered Trans-Tasman Patent Attorney & Systems Architect

---