

SAMPLE DELIVERABLE · AIG-01

AI System Register and Gap Assessment

The two artefacts the Australian Standards for AI will expect: the register template ready to fill, and a worked gap assessment showing what good looks like.

Register template + fictional worked sample · Red Cell Consulting

CONTENTS

01	Why now: the 2027 timetable
02	Your obligations, in plain language
03	The register: how to fill it
04	The register: template
05	Consequence rating and impact assessment
06	Gap assessment: worked sample
07	Remediation and the evidence file
08	About Red Cell Consulting

About this document

This is a sample of the deliverable Red Cell Consulting produces in a real engagement. Every organisation, figure and finding in it is fictional, and the structure, depth and standards are exactly what a client receives. Where a real engagement would cite source documents, the sample says so.

Why now: the 2027 timetable

Mandatory AI governance in Australia has a date. The organisations that will find it uneventful are the ones acting like it is already here.

THE TIMETABLE AS ANNOUNCED

WHEN	WHAT
July 2026	Australian Standards for AI announced; Office of AI established in the Department of the Prime Minister and Cabinet
August 2026	National Cabinet considers the design
Late 2026	Draft standards and implementation guidance expected
Early 2027	Legislation to Parliament
2027-2028	Staged enforcement

WHAT IS SIGNALLED FOR ORDINARY BUSINESSES

- A current register of the AI in use across the business and the data it touches
- Named accountable owners for AI systems
- Impact assessment for consequential uses, audit trails and incident reporting
- Consent rules around Australian creative content in training

The practical point

Registers take a quarter to build honestly. Remediation takes longer. Evidence, by definition, cannot be backdated: the file you start today is the file that answers 2027's questions. Everything in this pack is worth doing even if the final rules land narrower than signalled, because existing law already applies to your AI use.

Your obligations, in plain language

Before the Standards arrive, current law already governs AI use. The register you build for 2027 answers these questions too.

THE LAW THAT APPLIES TODAY

DUTY	HOW AI TRIGGERS IT	THE QUESTION YOUR REGISTER MUST ANSWER
Privacy	Personal information collected, inferred or used to train and prompt	Which systems touch personal information, and on what basis
Consumer law	Automated decisions, pricing, claims about your products	Where automated outputs reach customers, and who reviews them
Anti-discrimination	Screening, ranking or scoring people	Which systems affect individuals, and how bias is tested
Directors' duties	AI in material processes is a material risk	Whether the board can name its AI exposures and their owners
Work health and safety	AI-driven scheduling, monitoring and workloads	Whether machine-set demands are reviewed as work design

THE VOLUNTARY LAYER BECOMING EXPECTATION

The Voluntary AI Safety Standard, the six essential practices for AI adoption, and ISO/IEC 42001 all converge on the same grammar: know your systems, rate their consequences, name their owners, test and monitor, keep records. The Standards will harden that grammar; nothing in it will be wasted preparation.

Who this pack is for

Any organisation using AI at all, including AI that arrived embedded in ordinary software. If your business runs email, a CRM and an office suite, you have AI systems, and the register on the next pages has rows waiting.

The register: how to fill it

One honest afternoon starts it. The rules below keep it honest.

THE RULES

- **One row per system**, including embedded AI: office-suite assistants, CRM scoring, support-desk suggestions, transcription in your meeting tool. If it generates, ranks, predicts or decides, it rows.
- **Owner is a person**, not a team, and they know they own it. If nobody will own a system, that is your first finding.
- **Data honesty**: record what the system actually touches, not what the brochure says. Where you do not know, write unknown; unknown is a finding too.
- **Consequence before cleverness**: the rating (next page) asks what happens to a real person if the output is wrong, not how impressive the technology is.
- **Evidence links, even to gaps**: point to what exists (tests, approvals, supplier terms, incidents). Empty cells show the work; that is their job.

WHERE THE ROWS COME FROM

SOURCE	WHAT IT SURFACES
Finance: software subscriptions, 24 months	Sanctioned tools, incl. forgotten ones
IT: SSO and browser-extension inventory	Embedded and shadow usage
Team interviews, one hour per function	The workarounds people actually use
Supplier contracts	AI inside vendored services, and whose duty it is

Cadence

Quarterly review, thirty minutes, calendar-held: new rows, changed data flows, ratings still true, owners still owning. The review date is itself evidence.

The register: template

Two fictional example rows, then yours.

SYSTEM	WHAT IT DOES	DATA IT TOUCHES	SUPPLIER	OWNER	CONSEQUENCE	EVIDENCE
Example: Support reply assistant	Drafts customer responses for agent review	Customer names, order history	HelpDesk Co	J. Example	MEDIUM	Test set; June review
Example: CV screening aid	Ranks applications for shortlisting	Applicant PII	TalentTool	P. Example	HIGH	Bias test pending

CONSEQUENCE RATING · THE ONE QUESTION

RATING	WHAT HAPPENS TO A REAL PERSON IF THIS OUTPUT IS WRONG	MINIMUM CONTROLS
HIGH	Money, health, employment or legal position affected	Impact assessment, human decision point, tested before use, monitored, incident path
MEDIUM	Decisions influenced, but a person reviews before effect	Named reviewer, spot-testing, logged usage
LOW	Convenience only; errors are caught in normal work	Register row, owner, policy applies

Gap assessment: worked sample

A fictional organisation assessed against ISO/IEC 42001 domains and the signalled Standards trajectory. The pattern is what most organisations find.

CURRENT MATURITY VS REQUIRED · BY DOMAIN

DOMAIN	CURRENT	REQUIRED BY 2027	STATUS	PRIORITY
AI policy and leadership	78	70	MET	P3
Risk management	55	80	PARTIAL	P1
Impact assessment	30	75	GAP	P1
Data governance	62	80	PARTIAL	P2
Lifecycle controls	44	85	GAP	P1
Supplier management	38	70	GAP	P2
Monitoring and incident	58	85	PARTIAL	P2
Records and evidence	70	90	PARTIAL	P2

HOW SCORES ARE SET

Each domain is scored against written criteria (documented, practised, evidenced), by interview and by artefact inspection, never by self-assessment survey alone. "Current" means what could be shown to an outsider today; aspiration scores zero, which is the point of the exercise.

Reading the pattern

Leadership intent adequate, evidence and lifecycle controls short of 2027: this is the common shape. Less is missing than feared, and what is missing is exactly what a regulator asks for first, which is why the twelve-week sequence on the next page starts there.

Remediation and the evidence file

Twelve weeks from register to rehearsed. Each step exits with evidence, because the evidence is the deliverable.

THE SEQUENCE

WEEKS	WORK	EXIT EVIDENCE
1-2	Register completed; consequence ratings agreed; owners named and told	Signed register, review date scheduled
3-5	Impact assessments for the high-consequence systems	Assessments filed with decisions recorded
5-8	Lifecycle controls: testing gates and human sign-off wired into the two highest-risk workflows	Gate results in the file
8-10	Supplier terms reviewed against data and incident duties	Variance list with actions and dates
10-12	Monitoring live; incident runbook rehearsed; board briefed	Exercise record, board minute

THE EVIDENCE FILE · WHAT IT HOLDS

SECTION	CONTENTS
Register and reviews	Every version, every quarterly review note
Assessments	Impact assessments with the decision each produced
Testing	Test sets, results, dates, who ran them
Incidents	What happened, what was done, what changed
Suppliers	Terms reviews, variances, correspondence
Decisions	Board and executive minutes touching AI

One habit worth everything else

Date everything as it happens. In 2027, the creation dates in this file are themselves the proof that governance existed before it was demanded.

About Red Cell Consulting

Consulting and advisory, backed by platforms we design, build and operate ourselves.

Red Cell Consulting is a multi-disciplinary consultancy working across enterprise and technology, financial analysis, clinical psychology and behavioural science, advisory and investment readiness, restructuring, management consulting and AI governance. The distinctive part is that we also build: our platforms, including Recursys for deal intelligence, Myra for deterministic financial analysis and ADMRL for structured clinical reporting, run in production under the same disciplines we recommend to clients. When we advise on evidence, testing and governance, we are describing how our own systems work.

HOW AN ENGAGEMENT RUNS

STEP	WHAT HAPPENS
1 · Conversation	A free first discussion of the situation. We will tell you plainly whether we can help and what it would involve.
2 · Scope	A short written scope: the question, the deliverables, the timeline, the fee. Fixed where the work allows it.
3 · The work	Delivered in the structure this sample shows, with working sessions rather than big reveals. You see findings as they form.
4 · Handover	Deliverables that stand on their own, and the option of standing support where it makes sense.

FOR THIS SERVICE

AI governance at Red Cell is engineering first: we operate production AI under these disciplines ourselves, and our work builds the register, the assessments, the gates and the evidence file this pack describes. The rules arrive in 2027; the evidence should already exist. Start with the conversation.

Contact

+61 3 9448 4933 · Mon-Fri 9am-5pm Melbourne

info@redcc.au

www.redcc.au