AI GOVERNANCE · STEP TWO OF THREE

The AI risk assessment template

One use case, assessed properly, in about half a day. Built on the ISO/IEC 23894 process, with the National AI Centre's screening questions and the Australian contexts where the answer has to be no. Complete this before you write the policy, because the policy is written from it.

THE GAP THIS TEMPLATE CLOSES

Australian organisations that report implementing AI standards or guidelines **94%**

Australian organisations that assess third party vendor AI model claims **21%**

Fifth Quadrant for the National AI Centre, Australian Responsible AI Index 2025, published 26 August 2025.

Almost everybody has a document. One in five checks what their AI vendors actually claim. That distance is the difference between having a policy and having done the work.

WHO COMPLETES IT	TIME TO COMPLETE	COVERS	REVIEW AT
Business owner and risk	Half a day	One use case	12 months, or on change

METHOD

How to use this template

01

ISO/IEC 23894 mirrors ISO 31000 clause for clause. If your organisation already runs enterprise risk, you are extending a process you have, not building a new one.

1. **Scope, context and criteria**
   Part A. Name the specific use case, not "AI". Set the criteria for what counts as an unacceptable outcome before you look at any risks.

2. **Screen**
   Part B. Seven questions that set how deep the assessment needs to go, then the contexts where Australian law and guidance already draw the line.

3. **Identify**
   Part C. Work through the eight categories rather than brainstorming. Brainstorming finds the risks you already knew about.

4. **Analyse and evaluate**
   Part D. Rate likelihood and consequence, then compare the result against the criteria you set in Part A. Record inherent and residual risk separately.

5. **Treat**
   Part E. Four options and only four: avoid, mitigate, transfer, accept. Every accepted risk needs a named person accepting it.

6. **Monitor, review and record**
   Part F. Set the review date before you finish. The Commonwealth benchmark for high risk use cases is review at least every 12 months.

ONE USE CASE PER COPY

Assess "summarising client meeting notes in Copilot" or "drafting first pass tender responses". Do not assess "AI". A general assessment produces general controls, which is another way of saying no controls at all.

THE TEST OF A REAL ASSESSMENT

It has to be able to return the answer "not here". If the outcome was decided before you started, this is a procurement exercise wearing a risk assessment's clothes, and it will not survive contact with an auditor or an incident.

PART A

Define the use case

02

Use case name

Assessment date

What task does it do, in one sentence?
A person could read this and know exactly what is being assessed.

Tool, vendor and model

Who will use it

What data goes in?
List the classes. Personal information, client records, financial, contractual, health, classified.

Where is it processed and stored?
Country, and which sub-processors.

Is training on your inputs disabled?
Name the contract clause, not the marketing page.

What decision or output does it influence, and who is affected by it?

Where is the human review point?

Business owner, by name

Does a privacy impact assessment exist?
If personal information is involved and the answer is no, do that first.

Related assessments or references
PIA, security assessment, vendor review, register entry.

> **SET THE CRITERIA NOW, BEFORE YOU LOOK AT THE RISKS**
>
> Write down what would count as an unacceptable outcome for this use case, and what level of risk you are willing to carry. Doing this first is what stops the thresholds being reverse engineered from the answer you already want.

An unacceptable outcome for this use case would be

The level of risk we are prepared to carry here is

PART B

Screen the use case

03

Adapted from the National AI Centre's AI screening tool. A yes does not stop the use case. It sets how deep the rest of this assessment needs to go, and how senior the sign off has to be.

DOES THIS USE CASE...	YES	NO	UNSURE
Access personal, confidential or commercially sensitive information?	☐	☐	☐
Act autonomously, or at a scale where a person cannot realistically intervene?	☐	☐	☐
Affect people who are vulnerable, or who cannot opt out?	☐	☐	☐
Operate in a regulated area, or one with a professional duty attached?	☐	☐	☐
Produce harm that would be hard or slow to reverse?	☐	☐	☐
Influence a decision about a person's rights, entitlements, employment or money?	☐	☐	☐
Get used in a way the vendor's terms do not clearly permit?	☐	☐	☐

THE PROHIBITED CONTEXT CHECK

These are lines already drawn by Australian law, regulators or government policy. A yes in this table means the use case stops, or changes materially, before it proceeds.

WOULD THIS USE CASE INVOLVE...	YES	NO
Information security classified **OFFICIAL: Sensitive** or above in a public AI tool	☐	☐
Generating affidavits, witness statements, character references or other material tendered as evidence in NSW courts	☐	☐
Entering patient information into an AI tool without the patient's informed consent	☐	☐
Software used for a diagnostic or therapeutic purpose, without TGA consideration	☐	☐
A decision affecting a person's rights or entitlements, with no meaningful human review	☐	☐
Data your contracts, funding agreements or jurisdiction do not allow to leave the country	☐	☐

IF ANYTHING ABOVE IS TICKED YES

Stop. Redesign the use case, find a tool that can meet the constraint, or record a decision not to proceed. Do not carry it forward as a risk to be mitigated later.

Which one, and what changes as a result

PART C

Identify the risks

04

Work through every category, including the ones that look irrelevant. Write what could actually go wrong for *this* use case, not the category name again. Rate likelihood and consequence 1 to 5 using the scales in Part D.

REF	CATEGORY	WHAT COULD GO WRONG HERE	L	C	RATING
R1	**Privacy and personal information**				
R2	**Security, prompt injection and data exfiltration**				
R3	**Accuracy, and the cost of a confident wrong answer**				
R4	**Bias and fairness, and who is affected**				
R5	**Supply chain: vendor terms, sub-processors, jurisdiction**				
R6	**Identity and access, OAuth grants, offboarding**				
R7	**Legal and regulatory exposure in your sector**				
R8	**Explainability, and whether a decision can be contested**				
R9					
R10					

THE CATEGORY MOST OFTEN LEFT BLANK

R5. Only 21% of Australian organisations assess what their AI vendors claim about their models. Read the terms, name the sub-processors, and check what happens to your inputs. This is the row an auditor will turn to first.

PART D

Rate and tier

05

Likelihood

1. Rare. Would surprise everyone involved.
2. Unlikely, but has happened elsewhere.
3. Possible within a year of normal use.
4. Likely. Expect it at some point.
5. Near certain, or already happening.

Consequence

1. Negligible. Absorbed in normal work.
2. Minor. Internal rework, no external effect.
3. Moderate. A client, a complaint, a bad day.
4. Major. Notifiable breach, regulator, or real harm.
5. Severe. Serious harm, or the organisation's licence to operate.

RATING MATRIX

Consequence reads across, likelihood reads down.

LIKELIHOOD	C1	C2	C3	C4	C5
5 Near certain	Medium	High	Very high	Extreme	Extreme
4 Likely	Low	High	High	Very high	Extreme
3 Possible	Low	Medium	High	Very high	Very high
2 Unlikely	Low	Medium	Medium	High	Very high
1 Rare	Low	Low	Medium	High	High

WHAT EACH TIER REQUIRES

LOW	Business owner approves. Record the assessment and move on.
MEDIUM	Business owner plus risk. Named controls before go live, review at 12 months.
HIGH	Accountable officer approves. Controls in place and tested first, monitoring defined, review at 12 months or on any material change.
VERY HIGH	Executive or board approval, with the residual risk stated in writing and accepted by name.
EXTREME	Does not proceed in this form. Redesign the use case or record a decision not to proceed.

IF YOU ARE A NSW AGENCY, OR YOU SELL TO ONE

The NSW AI Assessment Framework is mandatory for all agencies and states that all generative AI solutions should be classified elevated risk. Where residual risk stays high or very high after mitigation, the use case must be referred to the NSW AI Review Committee.

PART E

Treat the risks

06

Four options, and only four. **Avoid** the risk by not doing it this way. **Mitigate** it with a control. **Transfer** it by contract or insurance. **Accept** it, by name. Anything that is not one of these four is a hope, not a treatment.

REF	OPTION	CONTROL OR ACTION	OWNER	IN PLACE BY	RESIDUAL

RECORD INHERENT AND RESIDUAL RISK SEPARATELY

The difference between the two is your control set. If they are the same number, either your controls do nothing or you have not been honest about the inherent rating. Both are worth finding out now rather than after an incident.

Highest inherent rating

Highest residual rating

Residual risk accepted, in plain words

What we are choosing to live with, and why that is reasonable.

PART F

The **decision record**

07

THIS USE CASE

- ☐ **Proceeds**
  Residual risk is within the criteria set in Part A.
- ☐ **Proceeds with conditions**
  Conditions recorded below, with owners and dates.
- ☐ **Does not proceed**
  Reason recorded below. This is a valid outcome.

Conditions, or the reason for not proceeding

ASSESSED BY	RISK OWNER	ACCOUNTABLE OFFICER	NEXT REVIEW DATE

NOW WRITE THE POLICY FROM THIS

WHAT THIS ASSESSMENT PRODUCED	THE POLICY CLAUSE IT BECOMES
Use cases assessed as acceptable	What AI may be used for
Data classes unacceptable at any tier	What must never be entered, named specifically
Vendors whose terms passed review	The approved tool list
Controls required to reach acceptable residual risk	Conditions of use, and who approves exceptions
Named risk owners	Accountability and escalation
Review cadence	When the policy is next revisited, and by whom

 Sources: ISO/IEC 23894:2023 · AS ISO/IEC 42001:2023 · National AI Centre, Guidance for AI Adoption and AI screening tool · DTA, AI impact assessment tool and staff guidance on public generative AI · Digital NSW, AI Assessment Framework · OAIC guidance on commercially available AI products · Supreme Court of NSW Practice Note SC Gen 23 · Ahpra · Fifth Quadrant for the National AI Centre, Australian Responsible AI Index 2025.