AUS GOV DATA SUMMIT • WORKSHOP B • APRIL 2026

Responsible AI in practice

This session raises questions that don't have tidy answers. The grey areas are where responsible use and judgement actually live, and that's both exciting and sobering. We made this pack to help guide your conversations beyond today's session.

CO-FACILITATORS

Siân Rinaldi & Britt Spyrou

April 2026

FOUNDATIONS

The three pillars of responsible AI

Responsible AI is not designed to slow down your delivery.
It is the guardrails that make confident delivery possible.

01

Enable opportunity.

Support agencies to use AI confidently, where it improves productivity, decisions and services. Reduce fear and fragmentation by setting clear expectations.

02

Build public trust.

Ensure government AI use is lawful, ethical and responsible. Apply risk-based, proportionate oversight where AI affects people or decisions. Make AI use transparent and accountable.

03

Adapt to change.

Recognise AI is evolving faster than legislation. Set requirements that scale with impact and maturity. Expect judgement, learning and iteration.

ACCOUNTABILITY

Organisation versus me

> You are not expected to have perfect answers today, but we need to ask the right questions so we can exercise judgement now.

The policy is phased, meaning mandatory requirements will be rolling in over time. Understanding both your agency's responsibilities and yours can help build confidence that you are acting appropriately and within the policy.

WHAT ORGANISATIONS CARRY

The system.

Structures and minimum settings.

- Clear accountability for AI use, with accountable officials and owned use cases.
- Visibility of in-scope AI use through internal registers.
- A public transparency statement about how AI is used and governed.
- A strategic approach to AI adoption and capability uplift.
- Formal processes to assess and manage higher-impact use cases.

WHAT YOU CARRY

The self.

Judgement and behaviour.

- Decide whether AI is appropriate for the task.
- Consider who could be affected and how.
- Use the rules you already know: privacy, data, records handling, accountability.
- Escalate when impact increases. You are not expected to carry the risk alone.

DECISION FLOW

The short-cut framework to the artificial intelligence impact assessment (AIIA) tool

Three stages. Use them to right-size governance to the actual impact of the use case.

01 Ethical design.

Design with ethics and existing guardrails in mind. Ask: is there a clear public benefit, is the data context appropriate, can existing controls manage this use?

02 Practical impact.

What effect does this actually have in the real world? Low, medium or high, based on who is affected, whether decisions or priorities are influenced, and scale or reuse.

> *Expect grey. A lot of this sits in judgement, not category.*

03 Proportionate governance.

What level of oversight is appropriate?

IMPACT LEVEL	WHAT GOVERNANCE LOOKS LIKE
Low impact	Existing controls, light touch.
Medium impact	Deliberate oversight, shared visibility.
High impact	Formal assessment, clear escalation and accountability.

REMINDER 01

Risk is driven by impact, not by technology sophistication.

REMINDER 02

If it changes who is prioritised or acted on, assume higher risk.

THE THINKING ENGINE

The Deficit-First Framework

AI sticks when four conditions are present. Treat them as a constraint set, not a checklist. All four need to be strong enough for the task. If one is weak, the rest will not carry it.

01 · WHY BOTHER

Deficit.

The reason anyone bothers. Start with a real friction point, capability gap, or task you avoid. Not a tool looking for a use. Deficits are personal. Two people in the same role will have different ones.

Test it: if the AI vanished tomorrow, would you fight to keep it? Automate the drudgery, protect the judgement, keep the fun or useful stuff you enjoy.

02 · WHERE IT SITS

Design.

Build into the work, not onto it. Map the full task: inputs, where AI sits in the sequence, what you still do, how the output gets used. If it feels like an extra step, the design is wrong.

At team level, co-design with the people who will use it. Plan different entry points for different deficits across the team. It's ok if AI use looks different across the team. I bet your inbox practices look different too.

03 · THE HABIT

Intent.

Driving the behaviour to get the value. Make AI use a deliberate habit, not a passing experiment. Build a reusable project, GPT, or template. Set the expectation that a given process now includes an AI step.

Without intent, you get a burst of enthusiasm. With it, capability compounds.

04 · THE FOUNDATION

Data.

Data readiness is structure, format, metadata, provenance, and context. Not just clean data. Favour tables and matrices over long documents and structure inputs deliberately. Work with your AI tool to understand what best feeds it useful information.

Poor outputs on day one does not build confidence to keep trying.

WHEN ADOPTION STALLS

When the framework breaks down

Four questions first. Four diagnostics second. Use this when adoption feels stuck and you need to find where.

01 Name one task in your role that is a genuine deficit for you. Why that one?

02 Where might AI help make it easier? What would you own? Where does AI sit?

03 Intentionally build the habit. If the behaviour happens, the value compounds.

04 What do your inputs need to look like on day one so the first result earns your trust?

IF THIS PILLAR IS WEAK	WHAT YOU WILL SEE, AND WHERE TO LOOK FIRST
Deficit	Someone tries the tool because it is new, not because they need it. Interest fades. Go back and name the real friction.
Design	Use stays ad hoc. AI sits next to the work instead of inside it. Redesign the task, not the prompt.
Intent	Enthusiasm at first, then drift back to old ways. Build the habit into rhythms, expectations, and measurement.
Data	Outputs are unreliable from the start. Fix inputs first: structure, format, metadata, context.

REFERENCE

Building a great prompt

Goal. Context. Source. Output. Each element does a different job. If any one is missing, the AI will guess. The guess is rarely what you needed.

ELEMENT	WHAT IT COVERS	EXAMPLE PHRASING
Goal	What you are trying to produce or achieve.	*"I have a 10-minute meeting with my Director and need talking points."*
Context	Who you are, your situation, what you are hoping to achieve.	*"I'm a policy officer in the department of X, preparing a brief for the minister on Y."*
Source	What documents or data the AI should draw from.	*"Reference the Australian Policy for the Responsible Use of AI. Be specific and ask me if you're not sure, don't guess."*
Output	The format, document type, length, and audience.	*"A three-sentence plain English summary, the three decisions being sought, two risks the FAS is likely to probe, and one piece of information that appears to be missing."*

WORKED EXAMPLE

A 40-page brief, and a meeting this afternoon.

You have read the document, but you don't feel prepared. A prompt that will actually help you:

```
Act as a senior policy advisor reviewing a brief. I have read the attached document but
need to ensure I haven't missed any critical details before briefing my Director this
afternoon. Please analyse this 40-page document and provide a concise briefing pack (max 2
pages equivalent) focusing on:

• Executive summary: the core proposal and the main decision required.
• Key risks: the top 3 high-level risks (political, financial, or implementation) the
  Director will probe on.
• Financial and legal implications: any significant funding requirements or legal hurdles.
• Stakeholder concerns: departments or stakeholders likely to object.
• Preparation questions: 5 tough questions the Director might ask me.

Please cite page numbers for all key points so I can locate them quickly.
```

WHERE TO FROM HERE?

Your next move

We know from the Copilot trial and a number of other research activities that AI is most likely to become entrenched if people jump in and try it. To build your maturity, have a think about what you could do tomorrow, next week and a month from now.

FOR LEADERS

What single governance move would add the most value at your current stage?

- Clarifying data ownership and reuse authority.
- Formalising stewardship and quality checks.
- Making acceptable and unacceptable AI use explicit.

FOR PRACTITIONERS

Where in your working week does something take too long that AI could support?

- Recurring tasks you set up manually each week.
- Documents you summarise or draft repeatedly.
- Briefing packs you condense for senior leaders.
- Meeting preparation and follow-up.

ONE RULE FOR BOTH

The best way to start is to pick one thing to try this week. See how you go, and consider how you could enhance your output. If it's just not sticking, try and understand why, but always keep trying something new.

STAY IN THE THINKING

Let’s keep thinking this through

Today raised questions that don't have tidy answers. The grey areas are where responsible use and judgement actually lives, which is both exciting and sobering.

If today sparked a question you want to keep working on, or you've got a situation you're wrestling with, we'd love to hear from you.

FACILITATOR

Siân Rinaldi

Helping people, organisations and communities understand and make best use of today's technology safely.

FACILITATOR

Britt Spyrou

Expert in the human architecture of systems: how we decide, deliver and adapt.