Data and AI governance

Setting yourself up for success

CO-FACILITATORS

Britt Spyrou & Siân Rinaldi

April 2026

HOUSEKEEPING

Welcome.

A quick orientation before we get into the work. Bathrooms, breaks, how today is structured, and what you should expect to walk out with.

WHAT TO EXPECT

Three blocks with activities. Plenty of time to talk. You will leave with a reference pack.

HOW TO PARTICIPATE

Your lived experience is the curriculum. Jump in, push back, share what is working and what is not.

LOGISTICS

Bathrooms to your left. Tea and coffee at the back. Breaks on the hour.

APS EXPERIENCE

Naming the elephants.

Three tensions shaping how the APS is thinking about AI right now. Worth naming before we design around them.

01

Generative versus agentic.

Government has been largely silent on agentic AI. The distinction matters for governance, risk, and how you scope what needs oversight.

02

Environmental trade-offs.

AI increases compute, but better data governance (fewer copies, cleaner structures) can offset waste. It’s not a clean one for one, but we can and should be thinking differently.

03

Pace versus readiness.

The AI Plan aims to substantially increase government use of AI to improve service delivery, policy outcomes and productivity. The pace is real. Readiness is uneven.

TODAY YOU WILL

Learning outcomes.

01

Understand AI.

Broad range of issues are on the table. Getting the basics right matters before anyone makes bigger moves.

02

Map your guardrails.

Understand what the Responsible Use of AI policy addresses. Know how to design use cases and apply governance effort proportionate to risk.

03

Take action.

Leave with a concrete, personal next step. A governance move for leaders, or a practical AI task for practitioners. Built on confidence, not fear.

01

FRAMING AND SHARED BASELINE

Our lived experience with AI.

Understanding the experience and skills in the room.

Give me one or two words to sum up how you are feeling about AI?

4 2 1 75

Give me one or two words to sum up how you are feeling about the APS using AI?

4 2 1 70

LIVED EXPERIENCE

Our lived experience.

AI, oh my.

Large language models (LLMs) generate stuff. OpenAI makes ChatGPT, Anthropic makes Claude, Microsoft makes Copilot (and note: Microsoft 365 Copilot is a different product).

Natural language processing.

Asking Google a question, you add keywords and Boolean operators (AND, OR, NOT, "xx").
Natural language processing lets you ask questions with more nuance.

> Where have you been most impressed or surprised by an AI output?
>
> Consider your own experience. Are you surprised by the polls?

LIVED EXPERIENCE

The Copilot trial: what they found.

7–8

Average prompt length in words.

1 in 3

Participants used Copilot daily.

75%

Of those who received three or more forms of training were confident using Copilot. 28% higher than those who received one form.

77%

Were optimistic about Microsoft 365 Copilot at the end of the trial.

70%+

Used Teams and Word during the trial, mainly for summarising and re-writing content.

~1 hr

Estimated efficiency gains per person, freeing time for higher value work.

86%

Of participants wanted to keep using Copilot after the trial ended.

Effort

Adoption required concerted effort and training to incorporate. It did not happen by default.

LIVED EXPERIENCE

Setting ourselves up for success.

Prompting is the way you engage with generative AI. The quality of what you get back depends on how you ask.

Scenario. You are going to the Aus Gov Data Summit next week, and you want to find out who is going to be there and how you should best prepare yourself. You send the team: "The Aus Gov Data Summit is next week. What do I need to know?"

Chad

Our enthusiastic grad.

Will give you everything, including things you did not ask for. Eager, wide, hard to focus. Not necessarily helpful.

Beryl

Well-seasoned policy officer.

Will clarify, narrow, and push back. Knows what matters and what does not. Knows who to ask, if she's not sure.

Grant

Your diligent Exec Officer.

Will structure the answer, cite sources, and flag what is still missing. Knows you, your priorities and contextualises.

LIVED EXPERIENCE

Prompting: your new superpower.

Four elements. Each does a different job. If any one is missing, the AI will guess. The guess is rarely what you needed.

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 am 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. If you do not tell it, it will guess."

Output.

Format, length, audience.

"Provide a three-sentence summary, three decisions being asked, two risks, and one thing that looks missing."

LIVED EXPERIENCE

The art of a good prompt.

A 40-page submission, a meeting this afternoon, and a need to be across every important bit without re-reading the whole thing.

Act as a senior policy advisor reviewing a submission. I have read the attached [document] but need to ensure I have not missed any critical details before briefing my Director this afternoon.

GOAL

Please analyse this 40-page document and provide a concise briefing pack (max 2 pages equivalent) focusing on the following:

CONTEXT

SOURCE

- Executive summary. What is the core proposal and the main decision required?
- Key risks. The top 3 high-level risks (political, financial, or implementation) that my Director will be worried about.
- Financial and legal implications. Significant funding requirements or legal hurdles.
- Stakeholder concerns. Which departments or stakeholders are likely to object?
- Preparation questions. Five tough questions the Band 1 might ask me, based on the text.

OUTPUT

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

LIVED EXPERIENCE

The art of a good prompt.

A 40-page submission, a meeting this afternoon, and a need to be across every important bit without re-reading the whole thing.

Act as a senior policy advisor reviewing a submission. I have read the attached [document] but need to ensure I have not missed any critical details before briefing my Director this afternoon. — GOAL / CONTEXT

Please analyse this 40-page document and provide a concise briefing pack (max 2 pages equivalent) focusing on the following: — SOURCE

- Executive summary. What is the core proposal and the main decision required?
- Key risks. The top 3 high-level risks (political, financial, or implementation) that my Director will be worried about.
- Financial and legal implications. Significant funding requirements or legal hurdles.
- Stakeholder concerns. Which departments or stakeholders are likely to object?
- Preparation questions. Five tough questions the Band 1 might ask me, based on the text.

— OUTPUT

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

02

MAP YOUR GUARDRAILS

What is responsible AI?

Exploring and applying the logic of the Responsible Use of AI policy.

RESPONSIBLE AI

Three pillars.

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.

JUDGEMENT AND BEHAVIOUR

Organisation versus me.

> You are not expected to have perfect answers today, but you are expected to exercise judgement now.
>
> The policy is phased. Mandatory requirements (formal assessments, owned registers, maturity uplift) will roll in over time.

ORGANISATIONS

The system.

- Clear accountability for AI use — accountable officials, 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.

INDIVIDUALS

The self.

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

RESPONSIBLE AI

The short-cut framework to the AIIA.

01

Ethical design.

Design with ethics and existing guardrails in mind. 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 and high impact is based on: who is affected, whether decisions or priorities are influenced, and scale and reuse. There are lots of grey areas.

03

Proportionate governance.

Match oversight to impact.

Low → existing controls, light touch

Medium → deliberate oversight, shared visibility

High → formal assessment, clear escalation

ACTIVITY

Use case: AI summaries.

An AI system is used by staff to draft summaries of briefing packs, committee papers, submissions or other long-form documents for SES and senior decision-makers.

Typical use:

- Condense large volumes of material into short summaries.
- Highlight key issues or themes.
- Generate first-draft executive summaries or talking points.

Summaries are drafts. A human author reviews, edits and approves the content before it is relied on or circulated.

The AI does not make decisions, prioritise actions, or determine outcomes.

The risk lives in: how visible human accountability remains, how summaries are treated in practice, and whether important nuance or caveats from source material are preserved.

Which existing controls already apply here?

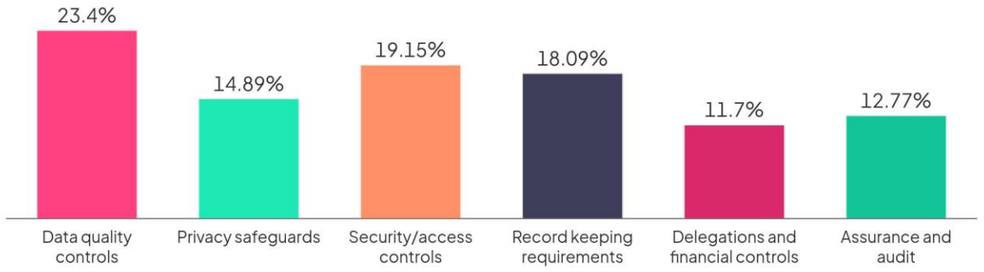

2 1 1 34

Practical impact

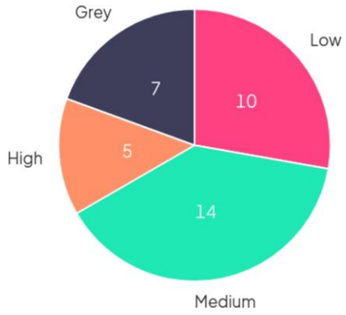

2 1 1 36

Which AI Ethics Principles are most relevant here?

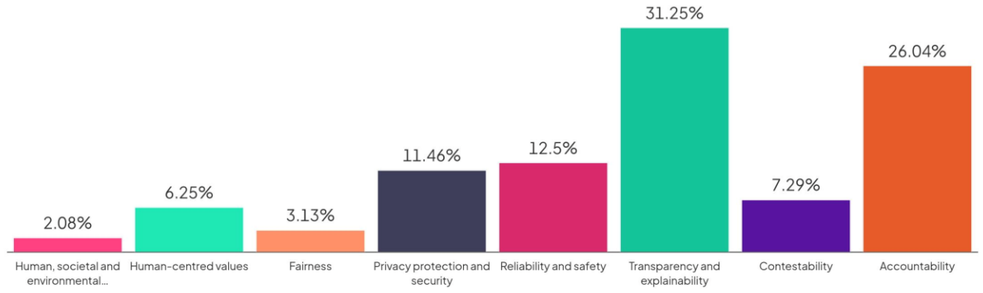

2 1 1 35

What are the proportionate governance controls needed ?

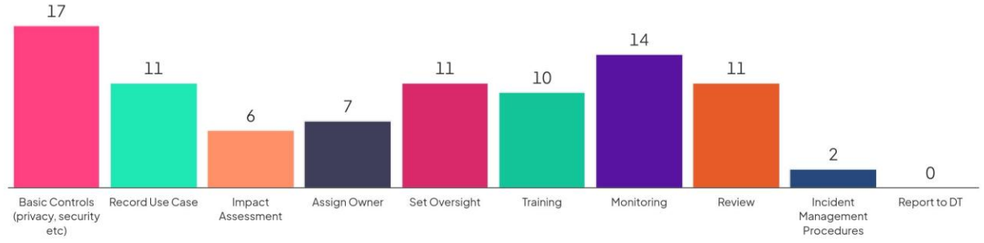

2 1 1 34

ACTIVITY

Use case: AI triage.

An AI system is used at the front end of an eligibility process for grants, programs or services, before a human assessor reviews applications.

Typical use:

- Screen applications based on patterns associated with likely eligibility.
- Flag "higher risk" or "lower priority" applications for additional scrutiny or delay.
- Prioritise the order in which applications are reviewed by humans.

Some applications are looked at sooner, some later, some with closer attention, based on AI-generated assessment rather than human judgement. Applicants may not know AI has influenced how quickly or carefully their case is handled.

The material question: although humans formally approve outcomes, AI shapes who progresses, when, and with what level of attention.

Which existing controls already apply here?

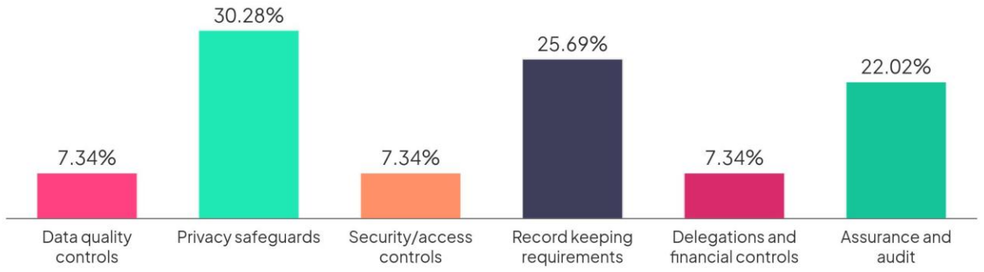

37

Practical impact

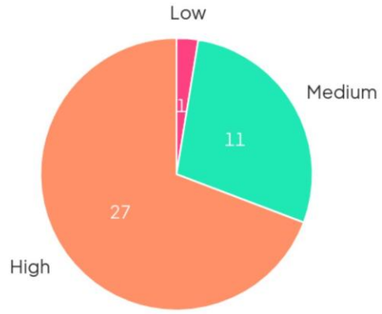

39

Which AI Ethics Principles are most relevant here?

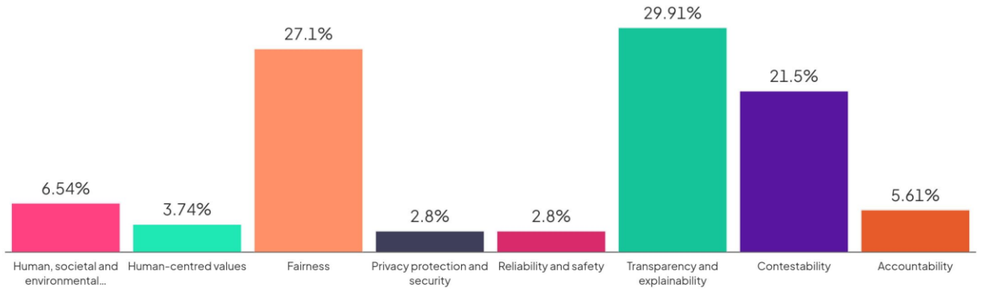

36

What are the proportionate governance controls needed ?

36

03

AI LEADERSHIP INTO THE FUTURE

Where to from here?

What can you take away from today to evolve your AI leadership?

COLLECTIVE PROGRESSION

Your maturity, your next move.

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

What's one part of your role where the time/value equation is unbalanced?

- 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.

INDIVIDUAL PROGRESSION

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.

If the AI vanished tomorrow, would you fight to keep it?

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. Rework the workflow, not the prompt.

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.

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

04 · THE FOUNDATION

Data.

Data readiness is structure, format, provenance and context. Not just clean data. Favour tables and matrices over long documents.

Poor outputs on day one are usually a data problem dressed up as an AI problem.

ACTIVITY

Think about your day.

01

Where have you seen people try AI processes and it not work out as planned?

02

How did the process stack up against the framework? Deficit, design, intent, data?

03

What is a part of work that you dread? How could you automate some or all of it? What treatment would your data need?

04

What will you do differently tomorrow?

AhaSlides

What was your key takeaway?

- AI might be useful with controls.
- Use AI to strengthen the documentation process
- 

- APS would benefit from a standardised AI/APS alignment benchmark set to help select appropriate models to use.
- It's complex
  There are good questions being asked
  We haven't really dealt with need vs want
  Enthusiasm is still trumping the inclsivity
- AI isn't going away, we need to continue to train our people and provide guidelines to enable them to find the best ways to use it.
- Checking unconscious bias
- Ros Kelly - Whiteboard
- I do know, what I don't know.
- Choose what your hate
- Don't lose critical thought and judgement from the process
- I need to invest in some reading glasses and/or I should have sat closer to the front.
- Considering using AI is an opportunity to review the controls we already have in place
- Gap in understanding of how to work with AI - collaborate, augmentation or agency.
- Controls important - but constant battle against uses of AI in inputs. (Echo chamber)
- Risk is driven by impact, not tooling
- How to apply real intelligence to artificial one whether it is controls, risks and governance etc.
- AI must be used with clear accountability
- Level of risk and impact defines the level and type of governance
- We need to use AI with proper controls and accountability
- Responsible AI use needs to take the human element into consideration.
- AI is a tool and we need to be able and confident in using it responsibly
- We need to ensure that AI is the best way to improve the workflow.
- AI is inevitable- moving forward how we can make the best out of it with good strategies, risk management and legislation system
- This is just three next step in technology evolution.
- We have a great AI Policy.
  But I can see a gap in AIIA, and a need in leadership awareness for accepting change in AI.

AhaSlides

What's something you're going to try ?

- Review the existing org controls in the context of AI
- Review existing ai policies and frameworks from APS beyond relevant state gov.
- Writing more specific requests for AI to process
- Be open minded but VERY critical of the 'how'
- To speak up when something doesn't sound right
- Try new areas where AI can assist..
  Improve awareness..
- Review actual use of AI against current organisation policies
- Think about the things I dislike and think about if AI can help
- Test out personal use of CoPilot more.
  Learn more, think more, use more.
- Get my own Chad to help tame my wild inbox.
- Put executive emails through an AI checker.
- Documentation
- Review the tasks or procedures that AI can apply, discussing with my team, engaging with stakeholders

Thank you !

Let’s keep the conversation going.

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.