



Australian Government



Test your scenario

AI risk and governance
tabletop exercises for teams

Facilitator guide

National Artificial Intelligence Centre

January 2026

Contents

Introduction	1
About the tabletop exercises	1
Role of the facilitator	2
Preparing for the session	2
Individual Scenarios	5
1. The helpful assistant	5
2. The Invisible Shortlist	6
3. The winning application	7

Introduction

AI risk and governance tabletop exercises for teams

First published: January 2026

Duration: 60 minutes

These tabletop exercises help business and organisational leaders explore what can go wrong when AI systems are introduced without clear governance, controls and oversight.

The scenarios are fictional but grounded in real-world situations.

They are designed to help your team:

- apply the National AI Centre's [essential AI practices](#)
- test how your business or organisation would respond in practice
- identify gaps in governance, policies and decision-making
- agree on practical actions to improve AI use in your business or organisation.

This facilitator guide explains how to run the exercises. Each activity is provided as a separate document and includes:

- a printable scenario briefing for participants
- discussion prompts
- a reporting template to capture insights and next steps.

You should read this guide alongside the essential AI practices before facilitating a session.

Participants should also review the essential AI practices before the exercise to ensure a shared baseline.

About the tabletop exercises

The exercises are designed to support guided discussion, not technical problem-solving. Participants do not need to be an AI or technology expert.

Each scenario focuses on common governance challenges, including:

- unclear accountability
- misleading or incorrect AI outputs
- workforce impacts and engagement
- supplier and procurement risk
- legal, regulatory and reputational exposure.

Using structured discussion and timed 'injects', teams explore how decisions are made under pressure and how well existing processes hold up.

All exercises can be run as stand-alone sessions, or combined for a longer workshop.

Available scenarios

These concepts have been generated with assistance from AI.

1. Sam, the helpful assistant

An AI chatbot improves efficiency but begins giving misleading information to customers.

2. The invisible shortlist

An AI recruitment screening tool raises concerns about fairness, transparency and discrimination.

3. The winning application

A grant application drafted with generative AI contains inaccurate citations and claims.

Role of the facilitator

The facilitator's role is to:

- guide discussion and keep the session on track
- introduce scenario updates at the right time
- encourage participation from all attendees
- keep the focus on governance and decision-making, not technical detail
- capture agreed actions, risks and gaps.

The facilitator does not need to be senior or a subject-matter expert. Anyone in the organisation can take on this role.

Preparing for the session

2 weeks before – preparation, participation and pre-reading

Planning

- Review the exercises and select the scenario(s) you will run.
- Confirm the session length (60 minutes is recommended per exercise).
- Decide whether the discussion will be in person or virtual.
- Book a suitable room or online meeting space.

You may adapt the discussion to suit your business or organisation, but do not change the core scenario content.

'Injects' are provided to increase pressure and move the discussion forward. Use them as written and introduce them progressively.

Participation

- Aim for 3 to 5 participants, including a nominated facilitator.
- The exercise can involve staff, executives, board members, or a mix.
- A diversity of roles and perspectives will strengthen the discussion.

Pre-reading

- Share the essential AI practices with all participants.
- Ensure participants understand that the session is about learning, not assessment.

On the day

Room setup and tone

- Set a no-blame environment
- Open the session by confirming that this is a safe space. The purpose is to identify risks and gaps in systems and processes – not to assign fault to individuals.

Materials

- Print copies of the participant scenario briefing and discussion guide.
- Keep facilitator notes and injects for facilitator use only.

Participants should receive the scenario for the first time on the day to simulate real-world conditions.

Suggested timing for each scenario (60 minutes total):

- Introduction and scenario briefing (10 minutes)
- Initial discussion (10 minutes)
- Scenario updates ('injects') and discussion (30 minutes)
- Wrap-up and agreed actions (10 minutes)

Timings are indicative. Adjust pacing to suit your group and discussion.

Facilitator Tips

Avoid rushing to solutions early. Some questions don't need a specific resolution. The aim is to surface trade-offs and gaps in governance, not to reach a single correct outcome.

- Ask for specifics. If someone says, "We would check the policy" ask:
 - Which policy?

- Who owns it?
- Where is it recorded?
- Avoid technical deep dives. Redirect discussion away from how the AI works and back to governance, accountability and decision-making.
- Note gaps as they arise. These exercises are designed to reveal missing processes, unclear ownership and untested assumptions.
- End with action. Assign practical next steps and clear accountability.
- Follow up. Circulate a short summary of decisions, risks and agreed actions after the session.

You do not need to ask every question on each slide. Use them as prompts to guide discussion. Select the questions that best surface decision-making, accountability and gaps in governance.

Individual Scenarios

1. The helpful assistant

Purpose

This scenario explores risks arising from customer-facing AI tools, including misinformation, accountability, disclosure and vendor reliance.

Facilitator focus

Display the [PowerPoint scenario](#) on screen and read the scenario to the group.

Work your way through the questions with the group.

Encourage the group to consider:

- who is accountable for AI decisions and outcomes
- how customer harm is identified and addressed
- how human oversight is maintained
- what controls apply to third-party tools.

Use the injects to steadily increase pressure and complexity.

Facilitation note

This scenario may prompt strong reactions, particularly where customer complaints, legal risk or reputational damage are raised. Reinforce the no-blame approach and redirect discussion toward organisational systems, governance arrangements and decision-making, rather than individual actions or technical failures.

Participants may focus on fixing the tool or blaming the vendor. Encourage the group to step back and consider accountability, oversight, escalation pathways and how customer harm is identified and addressed.

Bringing the discussion to a close

In the final discussion:

- avoid trying to resolve every issue raised
- focus on identifying **key risks**, **clear gaps**, and **practical next steps**.

Aim to capture:

- up to 3 priority governance issues

- one to 3 actions that could realistically be progressed in the next few months.

2. The Invisible Shortlist

Purpose

This scenario explores risks arising from the use of AI in recruitment, including fairness, accountability, transparency, workforce impact and reliance on vendor systems.

Facilitator focus

Display the [PowerPoint scenario](#) on screen and read the scenario to the group.

Work your way through the questions with the group.

Encourage the group to consider:

- who is accountable for recruitment decisions supported by AI
- how fairness and bias risks are identified and managed
- where human judgement and oversight sit in the process
- how staff, applicants and workforce representatives are engaged
- what controls and safeguards apply to vendor-provided tools.

Use the injects to steadily increase pressure and complexity.

Facilitation note

This scenario may raise sensitive issues, particularly for HR or people-management roles. Reinforce the no-blame approach and redirect discussion toward systems, governance and decision points, rather than individual fault.

Bringing the discussion to a close

In the final discussion:

- avoid trying to resolve every issue raised
- focus on identifying **key risks**, **clear gaps**, and **practical next steps**.

Aim to capture:

- up to 3 priority governance issues
- one to 3 actions that could realistically be progressed in the next few months.

3. The winning application

Purpose

This scenario explores risks arising from the use of generative AI in high-stakes organisational work, including accuracy, accountability, verification, integrity and disclosure.

Facilitator focus

Display the [PowerPoint scenario](#) on screen and read the scenario to the group.

Work your way through the questions with the group.

Encourage the group to consider:

- who is accountable for AI-assisted work and final outputs
- how accuracy and verification are managed under time pressure
- where human judgement and review processes apply
- how integrity, transparency and disclosure obligations are handled
- what governance expectations apply to boards, executives and staff.

Use the injects to steadily increase pressure and complexity.

Facilitation note

This scenario may raise sensitive issues, particularly where staff capability, workload pressure or organisational integrity are discussed. Reinforce the no-blame approach and redirect discussion toward systems, expectations, governance arrangements and review processes, rather than individual behaviour or good faith mistakes.

Participants may focus on intent or blame, or minimise the issue because outcomes were positive. Encourage the group to step back and consider accountability, verification requirements, disclosure decisions and how risks could be prevented in future.

Bringing the discussion to a close

In the final discussion:

- avoid trying to resolve every issue raised
- focus on identifying **key risks**, **clear gaps**, and **practical next steps**.

Aim to capture:

- up to 3 priority governance issues
- one to 3 actions that could realistically be progressed in the next few months.