



Artificial intelligence impact assessment tool

The Digital Transformation Agency (DTA) provides this impact assessment tool and its supporting guidance to assist Australian Government agencies to assess their proposed use of AI. Agencies should not treat the tool as legal advice or as authorising proposed AI use. Agencies are responsible for any decisions relating to their use of AI and for seeking technical and legal advice as appropriate.

For questions and suggestions on the impact assessment tool and guidance, please email ai@dta.gov.au

Digital Transformation Agency



© Commonwealth of Australia (Digital Transformation Agency) 2025

With the exception of the Commonwealth Coat of Arms and where otherwise noted, this product is provided under a Creative Commons Attribution 4.0 International Licence.

<http://creativecommons.org/licenses/by/4.0/legalcode>

The Digital Transformation Agency (DTA) has tried to make the information in this product as accurate as possible. However, it does not guarantee that the information is totally accurate or complete. Therefore, you should not solely rely on this information when making a commercial decision.

The DTA is committed to providing web accessible content wherever possible. If you are having difficulties accessing this document, please email ai@dta.gov.au.

Document control

This document is maintained by the DTA and supports implementation of the Australian Government Policy for the responsible use of AI in government (the AI policy).

It will be updated periodically as the policy and technology evolve. For enquiries, please email ai@dta.gov.au.

To view the AI policy, and to check you have the most up-to-date versions of the AI impact assessment tool and guidance, please visit <https://digital.gov.au/ai/ai-in-government-policy>.

Version	Description	Date
v1.0	Published	01/12/2025

Contents

- Introduction4
- 1. Basic information.....9
- 2. Purpose and expected benefits 14
- 3. Inherent risk assessment..... 16
- 4. Threshold assessment outcome.....32
- 5. Fairness34
- 6. Reliability and safety36
- 7. Privacy protection and security.....40
- 8. Transparency and explainability42
- 9. Contestability.....45
- 10. Human-centred values47
- 11. Accountability49
- 12. Use case review and next steps.....50
- Attachment A.....53

Introduction

The artificial intelligence (AI) impact assessment tool is for Australian Government teams working on an AI use case. It helps teams identify, assess and manage AI use case impacts and risks against [Australia's AI Ethics Principles](#). Understanding and managing AI use case impacts and risks is critical for effective AI governance and to fulfilling the Australian Government's commitment to safe and responsible use of AI.

The impact assessment tool and its supporting guidance are intended to complement and strengthen – not duplicate – existing frameworks, legislation and practices that relate to government AI use. It does this by focusing on AI-specific impacts and risks that existing approaches may not fully address. It does not replace a comprehensive risk management plan, which captures all risks, treatments and ongoing monitoring measures.

Policy for the responsible use of AI in government

On 1 December 2025, the Digital Transformation Agency (DTA) published an updated [Policy for the responsible use of AI in government](#) (the AI policy). The policy update strengthens government's approach to safe and responsible AI through new measures on AI governance. It includes a new mandatory requirement for agencies to conduct an AI impact assessment for use cases identified as in scope of the AI policy.

The updated AI policy provides implementation timeframes for agencies to meet the new requirements. Agencies are required to implement the AI impact assessment requirement for in-scope use cases by 15 December 2026. While agencies may need this time to action the new AI policy requirements, including mandatory impact assessment, agencies should implement them sooner if practicable.

Refer to the AI policy for more information on the definition of 'AI use case' and the AI impact assessment requirements.

Assessing officers should familiarise themselves with the AI policy and Australia's AI Ethics Principles. Also consider other DTA resources designed to support government AI adoption, including:

- the [AI technical standard](#)
- AI procurement resources, including the [Guidance on AI procurement in government](#), [AI contract template](#) and [AI model clauses](#)
- guidance on the use of public generative AI tools [for agencies](#) and [for staff](#).

The DTA piloted a previous draft of this tool – known as the ‘Pilot AI assurance framework’ – with volunteer agencies from September to November 2024 and published the pilot draft in October 2024. The pilot findings that shaped this version of the tool and guidance are outlined in the [Pilot implementation report](#). The title has been updated to ‘AI impact assessment tool’ to better reflect its intended scope and purpose.

The DTA welcomes user feedback on the tool and supporting guidance.

Please send questions or suggestions to ai@dta.gov.au.

AI impact assessment beyond the policy requirements

While the AI policy **requires** agencies to assess the impact of **in-scope** AI use cases, agencies are free to establish their own assessment requirements for AI use cases that are outside the scope of the AI policy.

Agencies may also use this tool to support other AI-related activities, such as procurements where suppliers use AI to provide goods and services. Refer to the [Guidance on AI procurement in government for further advice](#). Agencies should ensure specific internal requirements and expectations are clearly communicated to staff whose work involves AI.

For highly complex, novel or specific AI use cases, including AI-related activities which are outside the scope of the AI policy, agencies should also consider whether there are impacts which are not covered by this tool. These impacts must be assessed accordingly, and technical or legal advice sought where appropriate, as additional or different controls and mitigations may be required to address these impacts.

If you identify gaps in this tool, report these to the DTA. This feedback helps the DTA ensure that the AI policy and impact assessment tool remain fit for purpose and continue to evolve in response to emerging applications and risks.

Assessment roles and responsibilities

The tool and its supporting guidance are designed for Australian Government staff whose work involves AI. Each AI impact assessment must have an identified assessing officer and approving officer, and assessing officers should consult relevant experts for input. These roles are described below.

The updated AI policy also requires agencies to assign an **accountable use case owner** for each AI use case within the scope of the AI policy. This role is distinct from the assessment roles described above, although the accountable use case owner may also serve in one of those roles.

Assessing officer

An officer assigned to complete the assessment, coordinate the end-to-end process and serve as the contact point for any assessment queries. Depending on the specific use case and agency context, this officer may be a technical, data, governance or risk specialist, or they may be a policy or project officer from the business area that is implementing the AI use case in its operations.

Approving officer

An officer with appropriate authority to approve the AI use case assessment, including the inherent risk ratings. Like the assessing officer role above, the approving officer's specific role in the AI use case is not predetermined and will depend on the agency and use case context.

Expert contributors

Regardless of the assessing officer's role in the AI use case, they should seek peer review from colleagues and input from relevant experts as required, and document the experts consulted at section 1.10.

For some less complex, smaller scale AI use cases, the assessing officer may have all the information they need at their disposal to complete the assessment and may not require significant input from beyond their team. More complex use cases will likely require input from internal agency colleagues, including technical, data, risk management, policy and other domain area experts. If this expertise is not available in the agency, the assessing officer may need to seek external advice.

How to use this assessment tool

Check you have the latest version

The impact assessment tool will continue to evolve over time. To ensure you are using the latest version, please check <https://digital.gov.au/ai/ai-in-government-policy>.

Read the guidance

Before commencing this impact assessment tool and throughout the assessment process, you should [read the supporting guidance](#). The guidance mirrors the AI impact assessment tool's 12-section structure. For advice on completing a section in the tool, find the corresponding section number in the guidance.

Pre-assessment

Before assessing an AI use case, check if it is within the scope of the AI policy. The AI policy sets the criteria for in-scope use cases and the mandatory governance actions agencies must apply.

If your AI use case is out-of-scope of the AI policy, you may still find the impact assessment process useful. For example, impact assessment may support procurement, design and deployment decisions. Your agency may have specific requirements, in addition to the AI policy requirements, for assessing AI use case impacts. Check you are meeting any agency-specific requirements.

Start gathering use case information

You will need a broad range of information about the AI use case to complete the assessment. During the design phase the assessment is likely to be an iterative process, involving input from a range of experts. Regular updates may be needed as new information becomes available or design choices are refined.

Use case information you need for the assessment could include:

- the people the AI use case will affect, including demographic characteristics, needs or barriers they may face
- the potential impacts on individuals or groups, including direct and indirect impacts, their duration and reversibility, and how you intend to identify, assess and mitigate these
- the input data the AI system uses, including the type, source, collection method and security classification and how the input data will be stored and handled
- planned or existing security, monitoring, evaluation and quality assurance measures
- planned transparency measures to communicate information about the AI use case to individuals and groups
- how you will identify and consult relevant stakeholders
- how the AI system will record, log and explain any recommendations or decisions it makes, the extent of human oversight and validation of recommendations or decisions, and how any outcomes from the AI system will influence human decision-making
- who owns intellectual property rights in the inputs and outputs of the AI system, including copyright

- how your agency manages and delivers information technology services and solutions.

Threshold assessment: sections 1 to 4

First, complete sections 1 to 4 of the tool, which includes an assessment of inherent risks at section 3.

If all inherent risks are rated **low**, you can seek approving officer endorsement to conclude the impact assessment at section 4 and proceed with the use case, with appropriate plans for monitoring, evaluation and re-validation.

If any of the inherent risks at section 3 are rated **medium or high**, proceed to the full assessment and complete sections 5 to 12.

The AI policy also outlines additional requirements for use cases assessed as having an overall high inherent risk rating at the threshold assessment stage.

For any inherent risks rated medium or high, you may also consider seeking legal advice on whether the proposed use of AI is compliant with relevant laws and regulations before proceeding to the full assessment.

Full assessment: sections 5 to 12

The full impact assessment builds on the threshold risk assessment with more detailed analysis beyond the inherent risk level. It helps you examine specific potential harms, affected stakeholders, and contextual factors, to determine whether additional controls and mitigations are required.

Monitoring and evaluation

Regularly monitor and evaluate your AI use case throughout its lifecycle. If you identify a material change in the scope, usage or operation of the use case, you must formally re-validate your assessment, in line with AI policy requirements.

Re-validation

Check an approved use case assessment for accuracy and changes after deployment using re-validation. If re-validation results in assessment changes, the relevant officer or governance body must re-approve the changes. You may also specify other re-assessment intervals or triggers for your use case. This includes re-assessment to align with key project governance decision points.

1. Basic information

'AI use case' refers to the specific application of an AI system to achieve certain objectives or perform certain tasks. Refer to the AI policy for the definition of an AI use case and options to address situations where a single AI system encompasses multiple use cases.

1.1 AI use case profile

Name of AI use case

Enter your answer text here

Internal reference number or identifier

Enter your answer text here

Lead agency

Enter your answer text here

1.2 Establishing impact assessment responsibilities

Assessing officer

Officer responsible for completing this assessment, including coordinating input from relevant stakeholders, determining risk ratings and submitting the assessment for approval.

Assessing officer name

Enter your answer text here

Assessing officer position/team

Enter your answer text here

Assessing officer email

Enter your answer text here

Accountable use case owner(s)

If your use case is in scope of the AI policy, it must have a designated accountable use case owner. Refer to the AI policy and the Standard for accountability.

Name	Name	Position/team	Email

Approving officer

Approving officer name

Enter your answer text here

Approving officer position

Enter your answer text here

Approving officer email

Enter your answer text here

1.3 Additional roles and responsibilities

Record names and roles of people at your agency or third parties with responsibilities related to this AI use case. For example, responsibilities for developing the system, including external suppliers, monitoring the performance of the system or data governance.

Name	Role/title	Email	Responsibilities

1.4 AI use case description

In plain language, briefly explain how you are using or intend to use AI.

Enter your answer text here

1.5 In-scope use case

The following criteria will determine if the AI use case is in scope of the AI policy. If one of the below statements applies, your AI use case is in scope and must align with the requirements set out in the AI Policy. For details, refer to [Appendix C of the AI policy](#).

Select any criteria that place the use case within the scope of the AI policy.

- The use, misuse or failure of AI could lead to more than insignificant harm to individuals, communities, organisations, the environment or the collective rights of cultural groups including First Nations peoples.
- The use of AI will materially influence administrative decisions that affect individuals, communities, organisations, the environment or the collective rights of cultural groups including First Nations peoples.
- It is possible the public will directly interact with, or be significantly impacted by, the AI or its outputs without human review.
- The AI is designed to use personal or sensitive data, or security classified information.
- It is deemed an elevated risk AI use case as directed by the DTA.
- Not applicable.

1.6 Type of AI technology

Briefly explain the types of AI technology you are using or intend to use.

Enter your answer text here

1.7 Usage pattern

Select all that apply. Source: [Classification system for AI use](#).

- Decision-making and administrative action
- Analytics for insights
- Workplace productivity
- Image processing

Advisory note

Seek legal advice if AI automated decision-making is used for an administrative decision under an Act.

1.8 Administrative decisions

Only complete this section if you selected 'Decision-making and administrative action' in section 1.7 and if AI automated decision-making is used for an administrative decision under an Act.

Provide the legislative authority for automating such a decision.

Enter your answer text here

Advisory note

Express legislative authority is generally required to automate decision-making of an administrative decision under an Act. Legal advice should be obtained for any proposed use of AI in this context.

1.9 Domain

Select all that apply. Source: [Classification system for AI use](#)

- Service delivery
- Compliance and fraud detection
- Law enforcement, intelligence and security
- Policy and legal
- Scientific
- Corporate and enabling

1.10 Expert contributions

Record all experts you consulted during this impact assessment process, summarising their input and any resulting changes.

Name	Role and expertise	Date last consulted	Input and resulting changes

Advisory note

Agencies should consider engaging relevant internal or external expertise based on the complexity, novelty or potential impacts of the AI use case. This may include experts in ethics, law, discrimination, privacy, employment, intellectual property, technology, security, data, accessibility, automated decision-making, or the domain in which the AI system is being deployed.

1.11 Impact assessment review log

This impact assessment should be updated regularly throughout its lifecycle.

There are 2 types of review:

- a. **Pre-deployment review:** Continuous updates as part of the drafting process, before the assessment is approved. As new information emerges, design choices are refined and risks are reassessed, previous answers will be reviewed and updated.
- b. Post-deployment **re-validation:** Following the deployment of the use case, the AI policy requires agencies to:
 - regularly monitor and evaluate the use case to ensure it is operating as intended and that risks are effectively managed
 - re-validate the AI use case impact assessment by checking its accuracy and updating it when there is a material change in the scope, usage or operation of the use case.

Track impact assessment pre-deployment review and re-validation outcomes and summarise post-review changes.

Review type	Reasons for review	Review date	Post-review changes
<input type="checkbox"/> Pre-deployment			
<input type="checkbox"/> Re-validation			

2. Purpose and expected benefits

Under [*Australia's AI Ethics Principles*](#), the use of AI should have a clearly defined and beneficial purpose that is consistent with human, societal and environmental wellbeing.

2.1 Problem definition

Identify the problem you are trying to solve.

Enter your answer text here

2.2 AI use case purpose

Describe the purpose of your use of AI, focusing on how it will address the problem you have identified.

Enter your answer text here

2.3 Non-AI alternatives

Outline non-AI alternatives that could address this problem.

Enter your answer text here

2.4 Identifying stakeholders

Use the stakeholder mapping aid at **Attachment A** to identify stakeholder groups that may be affected by the AI use case. Record how they may be positively or negatively affected.

This will guide your consideration of expected benefits and potential risks. Stakeholder groups may include:

- end users of the AI outputs
- people who will be evaluated or monitored by the AI
- rights holders, for example copyright owners
- AI developers or engineers
- business or industry
- regulators
- agency staff and their unions, and other third-party personnel

- communities or other groups.

Stakeholder group	Briefly describe how they may be affected

2.5 Expected benefits

Consider the stakeholders identified in the previous question and briefly outline the expected benefits of the AI use case.

Consult the guidance for this section for resources to assist you.

What are the expected benefits of the AI use case?

Enter your answer text here

Advisory note

Your response should be supported by quantitative and/or qualitative analysis.

Qualitative analysis should consider whether there is an expected positive outcome and whether AI is a good fit to accomplish the relevant task, particularly compared to non-AI alternatives identified. Benefits may include gaining new insights or data.

3. Inherent risk assessment

This section identifies potential AI-specific risks that may arise from the AI use case. It is intended to complement existing risk management processes. It does not replace a comprehensive risk management plan capturing all risks, treatments and ongoing monitoring measures.

Base your impact assessment on the inherent risk, before it is treated with controls. Do not assess the residual risk, which would remain once controls are applied.

Consequence

Use the 5-tier consequence scale for consistent evaluation across impact areas. The consequence descriptions for sections 3.1 to 3.8 are summarised from the risk consequence assessment provided in the guidance. Refer to the table in the **guidance appendix** for further detail.

When assessing consequences, consider both intended and unintended consequences and outcomes. This includes evaluating the impact of system failure, the impact of misuse or malicious use, and of other deviations from expected use.

Likelihood

Use the risk likelihood scale in **Table 1** below to select a likelihood level for each of the impacts in sections 3.1 to 3.8.

Estimating likelihood for an AI use case can be harder due to limited historical data, changing behaviour and emerging challenges, such as the opaqueness of AI systems. If you are unsure, take a precautionary approach and rate the consequence as at least 'possible'. Seek advice from relevant experts – such as technical specialists, risk managers or domain experts – to inform or validate your assessment.

Risk rating

Use the risk matrix in **Table 2** below to determine your inherent risk rating for sections 3.1 to 3.8.

Provide a clear and concise rationale for each risk rating, explaining the relevant controls in place or planned. Ensure the assessment reflects the AI use case's scope, function, and current controls.

Table 1. Risk likelihood scale

Likelihood	Probability	Description
Almost certain	91% and above	Almost certain to eventuate within the foreseeable future
Likely	61-90%	Will probably eventuate within the foreseeable future
Possible	31-60%	May eventuate within the foreseeable future
Unlikely	5-30%	May eventuate at some time but is not likely to occur in the foreseeable future
Rare	Less than 5%	Will only eventuate in exceptional circumstances or as a result of a combination of unusual events

Table 2. Risk matrix

	Insignificant	Minor	Moderate	Major	Severe
Almost certain	Medium	Medium	High	High	High
Likely	Medium	Medium	Medium	High	High
Possible	Low	Medium	Medium	High	High
Unlikely	Low	Low	Medium	Medium	High
Rare	Low	Low	Low	Medium	Medium

3.1 Risk of reducing service accessibility and inclusion

This section asks you to assess the risk of your AI use making government services less accessible or inclusive for the public. Things to think about in assessing this risk include:

- Is the system being used for fundamental or critical government services to the public?
- Could any individuals or groups face new or increased barriers to accessing services or information?
- Could the system unintentionally exclude users with limited digital literacy or access to specific technologies?
- Could language barriers or cultural factors limit effective operation of the AI use case?
- Could the system produce answers which the public would misinterpret or find difficult to understand?

- Will the system allow users to seek verification of AI responses or review of AI outputs from a human? For example, will individuals be able to respond to misinformation produced by AI?

Why this matters

The use of AI can create barriers for some people if systems are not designed to be inclusive or accessible. This can range from minor inconvenience through to serious impacts.

Examples of minor inconvenience could include users not understanding how to interact with an AI system, or an AI system answering public enquiries with generic, inaccurate or confusing responses, requiring repeated attempts to resolve queries.

Serious impacts could include someone missing out on vital services because the AI has reached an incorrect decision without human input or oversight, or because the AI couldn't understand their language, disability or living situation. This could result in a breach of specific legislation relevant to the decision, if individuals are unable to access or understand the reasons for an AI-made decision or are unable to request a review of that decision.

What is the risk of your AI use making government services less accessible or inclusive for the public?

Consequence (select one)

- Insignificant – Minor glitch; no real barrier. Instantly resolved.
- Minor – Short-term, reversible barrier. Few users affected.
- Moderate – Noticeable access issues. Some groups impacted. Resolution requires effort.
- Major – Major service disruption. Widespread impact. Legal/public concern likely.
- Severe – Critical failure. Essential services inaccessible. Broad harm, urgent response.

Likelihood (select one)

- Almost certain
- Likely
- Possible
- Unlikely
- Rare

Risk rating (select one)

- Low
- Medium
- High

Rationale for risk rating including any existing controls

Enter your answer text here

3.2 Risk of unfair discrimination

This section asks you to assess the risk of your AI use unfairly discriminating against individuals, communities or groups. Things to think about in assessing this risk include:

- Could the system, through intended use, failure or misuse, replicate or amplify bias or discrimination? For example, from training data which is not diverse or representative, or embedded rules.
- Could certain groups be treated unfairly or inconsistently? This can happen if the system is trained on data reflecting existing inequalities or uses proxies correlating with race, class or gender and reproduces these biases in outputs. Proxies could include postcode, education level or country of birth.
- Is the AI system used as part of an assessment, eligibility, employment or resource allocation process? This includes making decisions with ethical or legal implications.
- Will users be able to identify bias in training data, and adjust the AI system accordingly? For example, users will gather more diverse sets of data if the AI system is reinforcing pre-existing patterns and therefore producing historical bias.

Why this matters

AI can produce discriminatory outcomes if:

- the parameters of the AI system are discriminatory
- the training data is unbalanced or contains hidden biases
- officials use AI outputs without proper scrutiny.

For example, if an AI system trained on biased hiring data screens job applications, it may disadvantage some demographics, causing issues from inconsistent shortlisting to excluding qualified candidates. This may not be compliant with Commonwealth anti-discrimination

legislation which prohibits discrimination on the basis of protected characteristics such as age, disability, race and sex.

What is the risk of your AI use unfairly discriminating against individuals, communities or groups?

Consequence (select one)

- Insignificant – Negligible discrimination; no harm. Issues caught and fixed early.
- Minor – Limited unfair treatment. Few individuals affected; quickly corrected.
- Moderate – Noticeable harm to some individuals/groups. Bias concerns raised; intervention needed.
- Major – Widespread harm to multiple communities. Significant impact; trust must be rebuilt.
- Severe – Systemic harm, especially to vulnerable groups. Public outrage; urgent reform required.

Likelihood (select one)

- Almost certain
- Likely
- Possible
- Unlikely
- Rare

Risk rating (select one)

- Low
- Medium
- High

Rationale for risk rating including any existing controls

Enter your answer text here

3.3 Risk of perpetuating stereotyping or demeaning representations

This section asks you to assess the risk of your AI use perpetuating stereotyping or demeaning representations of individuals, communities or groups. Things to think about in assessing this risk include:

- Could the training data cause the AI system to generalise the characteristics of individuals to a whole group? For example, not differentiating between individuals from racial or ethnic minorities.
- Could the system produce classifications or outputs that reinforce harmful stereotypes? For example, associating a certain role with a gender.
- Might outputs misrepresent individuals or communities, especially those historically marginalised? For example, the AI system processes and prioritises certain features of a particular race or ethnicity over others.
- Has the system been tested with diverse demographics, cultural contexts and perspectives?

Why this matters

AI systems can reinforce stereotypes or misrepresent people if they rely on overly broad categories or labels.

For example, if an AI tool used in social services classifies people based on postcode or income level, it might incorrectly assume certain behaviours or needs, leading to unfair treatment. This can result in minor impacts like poor communication, or more serious consequences such as inappropriate service responses. As noted above, this may not be compliant with Commonwealth anti-discrimination legislation.

Another example is if an AI tool publishes false information about an individual that causes significant harm to their reputation. This may bring the risk of defamatory action against an agency.

What is the risk of your AI use perpetuating stereotyping or demeaning representations of individuals, communities or groups?**Consequence (select one)**

- Insignificant – Mild stereotyping, quickly fixed. No lasting harm.
- Minor – Isolated incidents. Few affected. Promptly resolved.
- Moderate – Noticeable public concern. Some groups harmed. Needs intervention.
- Major – Widespread harmful stereotypes. Public outcry. Trust damaged.
- Severe – Systemic harm across communities. Legal risk. Urgent reforms needed.

Likelihood (select one)

- Almost certain
- Likely
- Possible
- Unlikely
- Rare

Risk rating (select one)

- Low
- Medium
- High

Rationale for risk rating including any existing controls

Enter your answer text here

3.4 Risk of harm

This section asks you to assess the risk of your AI use harming individuals, communities, groups, organisations or the environment. Things to think about in assessing this risk include:

- Could this specific system unintentionally cause physical, psychological, social, economic, reputational or environmental harm through its intended or unintended operation or outputs?
- Might this AI use case deprioritise or misclassify certain individuals or groups, resulting in harm or exclusion within the context in which it is deployed?

- Does the system interact with physical environments or control machinery as part of its defined use case?
- Is the AI system intended for use in the management or operation of critical infrastructure or services?

Why this matters

AI systems may cause direct or indirect harm through biased outcomes, exclusion, misinformation, poor decision-making, or unintended consequences.

For example, an AI system used to analyse drone footage for environmental compliance in a specific region may misidentify land clearing activities, leading to enforcement action against land holders without appropriate human oversight.

Assessment of environmental impacts should focus on the environmental consequences of the specific AI use case, not general or systemic environmental impacts of AI technology.

What is the risk of your AI use harming individuals, communities, groups, organisations or the environment?

Consequence (select one)

- Insignificant – Minor glitch. No real harm. Easily managed.
- Minor – Isolated public/business/environmental impact. Low cost. Resolved quickly.
- Moderate – Noticeable harm to people, businesses, or ecosystems. Public concern raised.
- Major – Widespread harm. Financial losses, distress, or ecosystem damage.
- Severe – Critical, irreversible harm. Economic and environmental crisis. Urgent action needed.

Likelihood (select one)

- Almost certain
- Likely
- Possible
- Unlikely
- Rare

Risk rating (select one)

- Low
- Medium
- High

Rationale for risk rating including any existing controls

Enter your answer text here

3.5 Risk of AI use raising privacy concerns

This section asks you to assess the risk of your AI use raising privacy concerns. Things to think about in assessing this risk include:

- Will the AI use case involve the collection, storage or disclosure of [personal information](#)? Does this include personal information that is [sensitive](#)?
- Does the use case involve a new or changed way of handling personal information?
- Is the AI use case likely to have a significant impact on the privacy of individuals?

Under the [Australian Government Agencies Privacy Code](#), for all 'high privacy risk' projects agencies must conduct a [privacy impact assessment \(PIA\)](#). To determine whether a PIA is required for your AI use case, you should complete a privacy threshold assessment (PTA). A PTA will help you identify your use case's potential privacy impacts and screen for factors that point to a 'high privacy risk', which requires a PIA. Use the outcomes of your PTA and, if required, your PIA, to inform your risk ratings below.

Why this matters

Use of AI is likely to involve new or changed ways of handling personal information that are likely to have a significant impact on the privacy of individuals. AI technologies rely on large data sets that often include personal information, which can create new specific privacy risks, amplify existing risks and lead to serious harms.

AI systems:

- often repurpose existing datasets for new functions or activities, which alters the purpose for using or disclosing personal information
- may require the integration of data from multiple sources, creating new pathways of collection, storage, or disclosure

- frequently rely on large-scale or sensitive datasets increasing the likelihood of high privacy risk – including biometric or behavioural data
- can sometimes enable re-identification or generate inferences that amount to new personal information, even where data is de-identified
- are known to generate inaccurate or false results increasing the risk that personal information collected, used and disclosed by an agency is inaccurate
- can produce personal information about an identified or reasonably identifiable individual that is inferred, incorrect or artificially generated.

You may also wish to consider the Office of the Australian Information Commissioner [Guidance on privacy and the use of commercially available AI products](#).

What is the risk of your AI use raising privacy concerns?

Consequence (select one)

- Insignificant – Minor data mishandling. No sensitive info exposed. Trust intact.
- Minor – Small data breach. Few affected. Resolved quickly.
- Moderate – Noticeable breach of sensitive data. Some distress caused.
- Major – Large-scale misuse of private data. Public trust damaged.
- Severe – Widespread sensitive data exposure. Severe harm and lasting loss of trust.

Likelihood (select one)

- Almost certain
- Likely
- Possible
- Unlikely
- Rare

Risk rating (select one)

- Low
- Medium
- High

Rationale for risk rating including any existing controls

Enter your answer text here

3.6 Risk raising security concerns – data aspects

This section asks you to assess the risk of your AI use raising security concerns due to the sensitivity or security classification of the data or information the AI system uses. Things to think about in assessing this risk include:

- Would a breach of this data cause operational, reputational, or personal harm?
- Is any of the data subject to legislative, government classification, contractual or equitable confidentiality obligations, or protected status? This can include:
 - protected information
 - security classified information
 - commercial information
 - law enforcement information
 - national security information
 - confidential information
 - personal information in accordance with the *Privacy Act 1988* (Cth).
- Are all external data sources integrated with the AI system under appropriate governance?
- What measures are in place to confine access and handling of data and information? For example:
 - Will inputs and outputs leave the agency's systems, and if yes, how will the data be handled. For example, how will it be transferred and stored?
 - Will any third-party operating the system have access to inputs and outputs, and can they use this data for other purposes. For example, can the system ingest data or conduct data mining?
 - What records will the system keep of access and use of data?
- What measures are in place to respond to unauthorised access, use or disclosure of sensitive or security classified data, such as a data breach or incident response plan?

Why this matters

AI systems handling sensitive or security classified data could pose significant security risks if not properly protected. Unauthorised access, data breaches or leaks, or misuse may:

- breach contractual or statutory obligations

- disrupt critical operations or services
- result in financial losses
- compromise national security and public safety
- expose individuals or systems to harm
- erode confidence in government data safeguards.

Agencies should have particular regard to the way the AI system will handle inputs and outputs. For example, commercially available AI may use inputs for further training and development of the AI or allow third parties to access outputs generated by the AI, through channels such as an interface with search engines. This may constitute a breach of contractual obligations not to disclose confidential information, or statutory obligations not to disclose protected information, such as obligations under federal taxation or health legislation.

What is the risk of your AI use raising security concerns due to the sensitivity or security classification of the data or information the AI system uses?

Consequence (select one)

- Insignificant – Minor lapse. No data misuse. Promptly fixed.
- Minor – Small breach. Few records accessed. Quickly secured.
- Moderate – Moderate data compromise. Privacy concerns raised. Needs investigation.
- Major – Major breach of sensitive data. Public trust shaken. Urgent response needed.
- Severe – Massive breach. National security and privacy at risk. Emergency overhaul required.

Likelihood (select one)

- Almost certain
- Likely
- Possible
- Unlikely
- Rare

Risk rating (select one)

- Low
- Medium
- High

Rationale for risk rating including any existing controls

Enter your answer text here

3.7 Risk of raising security concerns – system aspects

This section asks you to assess the risk of your AI use raising security concerns due to the implementation, sourcing or characteristics of the AI system? Things to think about in assessing this risk include:

- Is the AI system developed externally, or does it use open source or pretrained models?
- Can the system's behaviour be explained, tested, or audited?
- What safeguards does the AI system have if it malfunctions or produces harmful outputs? For example, can the agency interrupt or shut the system down via a circuit-breaker?
- Will the agency monitor, and take steps to address, signs of anomalies, dysfunctions and unexpected performance?
- Are there known vulnerabilities or dependencies, or likely security risks?
- Will the AI system be deployed in a secure environment on the agency's premises, or will it be deployed through the cloud? If deploying through the cloud, are the servers located in Australia or overseas?
- Is the AI system using novel or experimental algorithms without established security benchmarks?

Why this matters

The way an AI system is built, sourced or configured can introduce security vulnerabilities. These can range from low-impact issues – such as a brief service disruption due to minor misconfiguration – to a high-impact situation where a third-party component enables a major breach of sensitive data or system compromise.

It can also result in a breach of an agency's obligation to comply with the:

- information and resource protection requirements under the [Protective Security Policy Framework](#)
- requirements to protect personal information from misuse, interference and loss specified in the [Australian Privacy Principles](#).

It may also result in data being subject to foreign laws where that data is taken offshore.

What is the risk of your AI use raising security concerns due to the implementation, sourcing or characteristics of the AI system?

Consequence (select one)

- Insignificant – Minor flaw (e.g. bug). No real security impact. Quickly fixed.
- Minor – Small vulnerability exploited. Limited breach. Resolved with updates.
- Moderate – System feature causes data leak or access issue. Contained but serious.
- Major – System feature causes data leak or access issue. Contained but serious.
- Severe – Critical vulnerability causes widespread breach. National/public safety at risk. Emergency response needed.

Likelihood (select one)

- Almost certain
- Likely
- Possible
- Unlikely
- Rare

Risk rating (select one)

- Low
- Medium
- High

Rationale for risk rating including any existing controls

Enter your answer text here

3.8 Risk to reputation or public confidence

This section asks you to assess the risk of your AI use posing a reputational risk or undermining public confidence in the government? Things to think about in assessing this risk include:

- Is there a risk that the use case will fail public expectations of government as an exemplar?
- Even if technically sound, outcomes that appear unfair, insensitive, or misaligned with community values may generate media scrutiny and or erode public confidence.

Why this matters

Even where an AI system is ethical, technically sound and fully compliant, there may still be reputational risk arising from how it is used. Public perception can be shaped by the context, purpose, or outcomes of the system.

What is the risk of your AI use posing a reputational risk or undermining public confidence in the government?

Consequence (select one)

- Insignificant – Minor, isolated issue; negligible impact on public trust; quickly addressed.
- Minor – Brief concern or media attention; small reputational dent; resolved with prompt action.
- Moderate – Public questions oversight; moderate trust impact; requires remedial response.
- Major – Widespread criticism; major reputational damage; trust rebuilding needed.
- Severe – Profound loss of public trust; seen as governance failure; long-term recovery and reform required.

Likelihood (select one)

- Almost certain
- Likely
- Possible
- Unlikely
- Rare

Risk rating (select one)

- Low
- Medium
- High

Rationale for risk rating including any existing controls

Enter your answer text here

3.9 Overall inherent risk rating

Determine the overall inherent risk rating for the AI use case based on the ratings selected in sections 3.1 to 3.8 above. Use the highest risk rating identified in the sections above as the overall inherent risk rating. This ensures the assessment reflects the most significant potential risk exposure, even if the other risks are rated lower.

This overall inherent risk rating is the risk level under standard operating conditions, assuming only existing baseline controls are in place, before any additional controls are applied. It does not consider proposed mitigation strategies or enhancements that are yet to be implemented.

Overall inherent risk (select one)

- Low
- Medium
- High

Rationale for overall inherent risk rating and explanation of relevant risk controls

Enter your answer text here

4. Threshold assessment outcome

4.1 Assessing officer recommendation

If the assessing officer is satisfied that all risks in the inherent risk assessment at section 3.9 are **low**, they may recommend that a full assessment is not needed and that the agency accept the low risk and proceed with the AI use case.

If the overall inherent risk rating recorded at section 3.9 is **medium** or **high**, then you must complete a full assessment.

Assessing officer recommendation

- A full assessment **is** necessary for this use case.
- A full assessment **is not** necessary for this use case.

Comments (optional)

Enter your answer text here

Assessing officer name

Enter your answer text here

Assessing officer position

Enter your answer text here

Date

Enter your answer text here

4.2 Approving officer review

Approving officer decision

- I have reviewed the recommendation, am satisfied by the supporting analysis and agree that a full assessment **is** necessary for this use case.
- I have reviewed the recommendation, am satisfied by the supporting analysis and agree that a full assessment **is not** necessary for this use case.

Comments (optional)

Enter your answer text here

Approving officer name

Enter your answer text here

Approving officer position

Enter your answer text here

Date

Enter your answer text here

5. Fairness

This section helps agencies ensure that their AI use case aligns with [Australia's AI Ethics Principle](#) of fairness. The intention is to ensure agencies can clearly define fairness in operational terms and demonstrate that outcomes can be assessed for equity and impartiality.

Under the AI Ethics Principles, AI systems should, throughout their lifecycle, be inclusive and accessible and should not involve or result in unfair discrimination against individuals, communities or groups.

5.1 Defining fairness

Where appropriate, you should consult relevant domain experts, affected parties and stakeholders to determine how to define and contextualise fairness for your use of AI. Consider inclusion and accessibility, as well as discrimination and bias. Consult the guidance for prompts and resources to assist you.

Do you have a clear definition of what constitutes a fair process and/or outcome in the context of your use of AI?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

5.2 Measuring fairness

Measuring fairness is an important step in identifying and mitigating fairness risks. A wide range of metrics are available to address various concepts of fairness. Consult the guidance for resources to assist you.

Do you have a way of measuring the fairness of system outcomes quantitatively or qualitatively?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

6. Reliability and safety

This section helps agencies ensure that their AI use case aligns with [Australia's AI Ethics Principle](#) of reliability and safety. Its intention is to ensure agencies consider reliability and safety in operational terms and demonstrate how these will be measured, monitored and acted on.

Under the AI Ethics Principles, AI systems should throughout their lifecycle reliably operate in accordance with their intended purpose.

6.1 Data suitability

Consider data quality and factors such as accuracy, timeliness, completeness, consistency, lineage, provenance and volume. Also ensure that all necessary data permissions and rights are in place.

If your AI system requires the input of data to operate, or you are training or evaluating an AI model, can you explain why the chosen data is suitable for your use case?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

6.2 Indigenous data

Consider if use of Indigenous data and AI outputs is consistent with the expectations of First Nations people, and the [Framework for Governance of Indigenous Data \(GID\)](#). See further advice in the guidance.

If your AI system uses Indigenous data, including where any outputs relate to Indigenous people, have you ensured that your AI use case is consistent with the Framework for Governance of Indigenous Data?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

6.3 Suitability of procured AI system

This may include the procurement of a model, system, multiple systems or a class of systems from a third-party provider. It also covers situations where you are using open-source systems, application programming interfaces (APIs) or otherwise sourcing or adapting existing systems. Factors to consider are outlined in the guidance.

This includes ensuring that the AI is suitable for your intended use and that contracts used to procure the AI model or system protect against relevant risks. Ensure there are appropriate terms covering factors such as:

- testing
- support
- data governance
- privacy
- security
- intellectual property provisions
- warranties.

If you are procuring an AI model, can you explain its suitability for your use case?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

6.4 Testing

Outline any areas of concern in results from testing. If the AI system has not yet undergone testing, outline elements to be considered in testing plan. For example, the model's accuracy.

Has the AI system been tested sufficiently and are you satisfied with its reliability and safety for the context of your use case?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

6.5 Pilot

If answering 'yes', explain what you have learned or hope to learn in relation to reliability and safety and, if applicable, outline how you adjusted the use of AI.

Have you conducted, or will you conduct, a pilot of your use case before deploying?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

6.6 Monitoring

Explain how you will monitor and evaluate performance.

Have you established a plan to monitor and evaluate the performance of your AI system?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

6.7 Preparedness to intervene or disengage

Refer to the guidance for resources to assist you in establishing appropriate processes.

Have you established clear processes for human intervention or safely disengaging the AI system where necessary? For example, if stakeholders raise valid concerns with insights or decisions or an unresolvable issue is identified.

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

6.8 Training of AI system operators

An operator of the system may include anyone who uses the AI to support their work, such as staff interpreting insights, responding to system alerts, or making decisions informed by the AI's outputs. With all automated systems, including AI systems, there is always the risk of overreliance on results. Operators should receive appropriate training to understand how to use the system responsibly, monitor and critically evaluate its outcomes.

Is there a process in place to ensure operators of the AI system are sufficiently skilled and trained?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

7. Privacy protection and security

This section helps agencies ensure their AI use case upholds [Australia's AI Ethics Principle](#) of privacy protection and security by ensuring that personal information is collected, used, stored and shared in a way that minimises risk and protects individuals' privacy in accordance with the [Australian Privacy Principles \(APPs\)](#).

Agencies must ensure their AI use case operates within an appropriate security framework, including the Protective Security Policy Framework (PSPF), with safeguards proportionate to the sensitivity of the data and potential harm.

Under the AI Ethics Principles, AI systems should throughout their lifecycle respect and uphold privacy rights and data protection, and ensure data security.

7.1 Minimise and protect personal information

See the guidance for advice on:

- collecting, using and disclosing personal information in accordance with the APPs
- privacy enhancing technologies.

You should also consider the Office of the Australian Information Commissioner [Guidance on AI and Privacy](#).

Are you satisfied that any collection, use or disclosure of personal information complies with the APPs?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

7.2 Privacy threshold and/or impact assessment

Record where you have stored your privacy threshold assessment and, if applicable, your privacy impact assessment are stored. For example, the document management system, case file, or other repository.

Enter your answer text here

7.3 Security risks

Engage with your agency's IT Security Adviser (ITSA) and consider the latest security guidance and strategies for AI use, such as the:

- Protective Security Policy Framework (PSPF)
- Information Security Manual (ISM)
- Australian Signals Directorate advice on [Engaging with AI](#).

What measures are in place to address security risks arising from the operation of the AI?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

8. Transparency and explainability

This section supports agencies to uphold [Australia's AI Ethics Principles](#) of transparency and explainability, which entail:

- consulting with affected stakeholders
- maintaining clear and accurate records
- disclosing when AI is used
- providing appropriate explanations of AI outputs that allow individuals to understand, question, or seek review where appropriate.

Under the AI Ethics Principles, there should be transparency and responsible disclosure so people can understand when they are being significantly impacted by AI and can find out when an AI system is engaging with them.

8.1 Consultation

Refer to the list of stakeholders identified in section 2.4. Seek out representatives with the appropriate skills, knowledge or experience to engage with AI ethics issues. Consult the guidance for prompts and resources to assist you.

Have you consulted stakeholders identified in section 2.4?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

8.2 Public visibility

When deciding to publish or not publish AI use information, consider the advice in the guidance on appropriate transparency mechanisms, information to include and factors to consider.

Will appropriate information be made publicly available, such as the scope and goals related to the use of AI?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

8.3 Maintain appropriate documentation and records

Comply with documentation and record-keeping requirements, including maintaining reliable records of decisions, testing and information and data assets used in an AI system. This is important to enable internal and external scrutiny, continuity of knowledge and accountability.

Have you put in place processes to maintain appropriate documentation and records throughout the lifecycle of the AI use case?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

8.4 Disclosing AI interactions and outputs

Consider members of the public, agency staff or other stakeholders who may interact with the AI system, or decision-makers who may rely on its outputs.

Will people be informed when they interact with the AI system or receive outputs generated by AI?

- Yes
- No
- Not applicable

Explain your answer, including, if applicable, how AI involvement will be disclosed.

Enter your answer text here

Will members of the public interacting with the system be provided with the ability to request a non-AI alternative?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

8.5 Offer appropriate explanations

You should be able to clearly explain how a government decision or outcome has been made or informed by AI. These explanations should be understandable to both technical and non-technical audiences. This is especially important if your AI system will materially influence administrative action or decision-making about individuals, groups, organisations or communities.

Will your AI system allow for appropriate explanation of the factors leading to AI-generated decisions, recommendations or insights?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

9. Contestability

This section supports agencies to uphold [Australia's AI Ethics Principle](#) of contestability. Agencies need to have clear processes for notifying individuals about AI use in administrative actions and for managing review or appeals in a timely manner.

Under the AI Ethics Principles, when an AI system significantly impacts a person, community, group or environment, there should be a timely process to allow people to challenge the use or outcomes of the AI system.

9.1 Notification of AI affecting rights

See the guidance for help interpreting 'administrative action', 'materially influenced' and 'legal or significant effect' as well as recommendations for notification content.

Will individuals, groups, organisations or communities be notified if an administrative action with a legal or significant effect on them was materially influenced by the AI system?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

9.2 Challenging administrative actions influenced by AI

Administrative law is the body of law that regulates government administrative action. Access to review of government administrative action is a key component of access to justice. Consistent with best practice in administrative action, ensure that no person could lose a right, privilege or entitlement without access to a review process or an effective way to challenge an AI-generated or informed decision. The use of AI should not make it more difficult to access administrative law rights or place the agency in a position where it might contravene its administrative law duties.

Is there a timely and accessible process to challenge the administrative actions discussed at section 9.1?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

10. Human-centred values

Under [Australia's AI Ethics Principles](#), AI systems should throughout their lifecycle respect human rights, diversity and the autonomy of individuals.

This section supports agencies to uphold the AI Ethics Principle of human-centred values. It is intended to help agencies consider AI use case alignments with Australia's human rights obligations and to incorporate diverse perspectives. Incorporating diverse perspectives can help avoid biased or discriminatory outcomes that could undermine individual or community well-being.

10.1 Incorporating diversity

Consider how you have incorporated diverse perspectives through the lifecycle of your AI use case. For example, consider the choice of data, composition of development and deployment teams and the stakeholder and user groups to choose to consult.

Are you satisfied that you have incorporated diversity and people with appropriately diverse skills, experiences and backgrounds throughout the lifecycle of your AI use case?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

10.2 Human rights obligations

Complete this question after completing previous sections of the assessment. This approach will enable a more considered assessment of the human rights implications of your AI use case.

Have you consulted an appropriate source of advice or otherwise ensured that your AI use case and the use of data align with human rights obligations?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

11. Accountability

Under [Australia's AI Ethics Principles](#), AI systems should throughout their lifecycle be identifiable and accountable for the outcomes of the AI systems, and human oversight of AI systems should be enabled.

This section supports agencies to uphold the AI Ethics Principle of accountability. It is intended to help agencies consider its responsibility for the outcomes of the AI systems that they design, develop, deploy and operate.

11.1 Ensuring accountability during the lifecycle of the AI system

Consider the mechanisms you should put in place to ensure accountability for AI systems and their outcomes. This includes both before and after their design, development, deployment and operation.

Are there mechanisms in place to ensure responsibility and accountability by the agency and relevant individuals for AI systems and their outcomes?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

12. Use case review and next steps

12.1 Alignment with relevant legal frameworks

In addition to broader legislative obligations such as security, information, privacy, discrimination and human rights considered in sections 5 to 11, agencies must consider alignment with legislation and regulatory instruments specific to their context.

Have you identified and documented all agency specific legislation, or regulatory instruments that are relevant to the AI use case?

- Yes
- No
- Not applicable

Explain your answer

Enter your answer text here

12.2 Legal advice

Identifying and appropriately documenting the need for legal advice is an important part of managing risk and impacts. This section asks you to confirm if legal advice was sought during the assessment process and where that advice is recorded.

This information should not be disclosed to anyone other than those who need to know or access the information within the agency.

Have you identified the need for legal advice at any stage during the assessment process?

- Yes
- No
- Not applicable

If yes, where is the legal advice stored? For example, note the document management system, case file or other repository:

Enter your answer text here

12.3 Risk summary table

In the table below, list any inherent risks rated as **medium** or **high** in section 3 of the threshold assessment. Briefly explain the mitigations or controls that have been or will be applied and how these mitigations have influenced the residual risk rating.

Risk title	Risk treatments	Residual risk
<i>Enter risk title</i>	<i>Explain risk treatments</i>	<input type="checkbox"/> Low <input type="checkbox"/> Medium <input type="checkbox"/> High

12.4 Overall residual risk rating

Based on the assessment of residual risk in section 12.3, choose an overall residual risk rating for the AI use case.

Residual risk rating (select one)

- Low
- Medium
- High

Rationale for overall residual risk rating and explanation of relevant risk controls

Enter your answer text here

12.5 Internal governance body review

If your use case's inherent risk is rated **high** at section 3, you are required under the updated AI policy to apply specific actions, including submitting it to an appropriate internal governance body for review. Refer to the AI policy for further information on this requirement.

Your agency may also have specific requirements for referring AI use cases to an internal governance body, separate from the AI policy. Refer to your agency's AI-specific policies for more information.

Record the outcome of internal governance body review in the table below, including any recommendations and agreed next steps, if applicable.

Internal agency governance body recommendation

Enter your answer text here

Attachment A: Stakeholder mapping aid

Assessment section 2.4 asks you to identify stakeholder groups that may be affected by your AI use case, including its underlying AI system(s). The table below may assist with this exercise.

Please note the table has been provided as a discussion prompt and is not intended as a prescriptive or comprehensive list.

Stakeholder type	Identified use case stakeholders	Potential impacts (positive or negative)
<p>End users</p> <p>People who will use the AI system and / or interpret its outputs.</p>		
<p>Evaluation or decision subjects</p> <p>People or groups who will be evaluated or monitored by the AI system (e.g. who the system is making predictions or recommendations about).</p>		
<p>Oversight team</p> <p>The person or team who is managing, operating, overseeing or controlling and monitoring the AI system throughout its lifecycle, including information managers.</p>		
<p>System owner or deployer</p> <p>The executive or executives responsible for authorising particular uses of an AI system.</p>		
<p>AI model or AI system engineers</p> <p>Those involved in AI model or system design, development and maintenance.</p>		

Stakeholder type	Identified use case stakeholders	Potential impacts (positive or negative)
<p>Rights holders</p> <p>Those who hold the rights to materials used by this AI use case (e.g. copyright owners or creators)</p>		
<p>Malicious actors</p> <p>Those who may intentionally misuse the AI system.</p>		
<p>Bystanders</p> <p>People who may be affected by the AI system without actively engaging with it, including those in the system's physical operating environment.</p>		
<p>Regulators</p> <p>Authorities responsible for creating, enforcing, or monitoring compliance with laws, regulations, and standards relevant to the AI system or its use.</p>		
<p>Other government agencies</p> <p>Agencies beyond your own agency that may be impacted by, collaborate on, or have policy or oversight responsibilities related to the AI system and its use.</p>		
<p>Civil society organisations</p> <p>Non-governmental organisations, advocacy groups, and community representatives concerned with ethical, social, or rights-based implications of the AI system and its use.</p>		

Stakeholder type	Identified use case stakeholders	Potential impacts (positive or negative)
<p>Communities or groups</p> <p>Communities or groups that are likely to be affected by the AI system or its use, including those that may be vulnerable or disproportionately impacted.</p>		
<p>Associated parties</p> <p>Third parties indirectly impacted by the AI system’s evaluations or decisions, as well as those with an interest arising from their connection to other stakeholders—such as individuals, organisations, businesses, or industries.</p>		
<p>Staff and other personnel</p> <p>Personnel whose roles and workflows may be affected by your AI use case, including agency staff and their unions, as well as contractors and other third-party personnel.</p>		
<p>Intermediaries</p> <p>A facilitator or agent between 2 parties whose role may evolve with AI integration (e.g. tax agents).</p>		