Attachment D:
POLICY– Use of Artificial Intelligence (AI)

About this document	
Purpose	This Policy sets out the principles the Office follows for the ethical, safe and responsible use of Artificial Intelligence (AI) by staff.
User/s	All staff of the Office of the Commonwealth Ombudsman (the Office).
Publication/release to other sites	This Policy will be released to the VOLT platform, which is internal to the Office, and published to the Ombudsman website.
Outcome	Use of AI in the Office is: • Ethical, safe and responsible. • Consistent with legal and regulatory responsibilities
Version number	2.0
Consultation	AI Principal Advisors
Approved/endorsed by	Executive Committee
Date approved/endorsed	24 February 2026
Next review date	24 March 2027
Contact team	ICT & Security Team, Corporate Branch

Contents

Purpose

The Use of Artificial Intelligence Policy (the policy) specifies how the Office of the Commonwealth Ombudsman (the Office) will use Artificial Intelligence (AI) technology and systems. The policy sets out the general principles the Office follows to ensure any use of AI technology in our work is ethical, safe, responsible and in line with community expectations and administrative law principles.

Specifically, this policy aims to:

- protect the rights of stakeholders, including the Australian community
- strengthen public trust in the Office's use of AI by providing enhanced transparency, governance and risk assurance
- engage and empower our staff to use AI responsibly
- embed an adaptable approach that can evolve and develop over time
- implement the requirements of the Digital Transformation Agency ("DTA")'s Policy for the responsible use of AI in government v 2.0. ("DTA AI Policy")

Scope

This Policy applies to all Office staff, including ongoing, non-ongoing and casual employees. It also applies to contractors, consultants and visitors who have access to the Office's ICT equipment, systems and services.

This Policy applies to all AI systems under the Office's control, including those developed in-house, purchased from vendors or embedded within larger software platforms including cloud-based systems.

Relationship to other policy documents

This Policy is intended to be consistent with the Office's ICT security policies, and whole of government policy on the use of AI in government. Any obligations imposed by this policy are in addition to:

- Existing controls on the use of technology including the Policy for the responsible use of AI in government (Version 2.0), December 2025, the AI Acceptable Use Policy, and the ICT Security Policy
- Existing ethical and conduct obligations including the APS Code of Conduct

- Existing obligations regarding the treatment of sensitive information including the Privacy Act, the Office Privacy Policy, and the Protective Security Policy Framework (PSPF).
- Existing obligations regarding procurement including the Public Governance, Performance and Accountability Act 2013, the Office Procurement Policy and the DTA's guidance on procurement of AI technologies.

What is AI?

An AI System is any technology that uses data to make inferences and generate outputs such as predictions, recommendations, or decisions with a degree of autonomy.

This includes, but is not limited to:

- machine learning models
- generative AI tools (such as Chat GPT)
- predictive analytic systems
- chatbots that generate their own responses.

It excludes:

- formulae and calculations that produce a fixed, predictable output, like a spreadsheet
- rule-based automations
- traditional business intelligence dashboards.

This definition is consistent with the definition that has been adopted by the Organisation for Economic Co-operation and Development (OECD), and the DTA Policy. If you are unsure about whether a technology falls under this policy, consult the AI Accountable Officials.

AI Use Cases

This policy uses the term 'AI Use Case' consistent with the DTA Policy - An AI use case is a specific application of an AI system or systems to achieve certain objectives or perform certain tasks. Any AI use by staff must have an approved use case evaluation in place before use.

A single AI system can have one or more applications, which can each differ in their intended purpose, functionality and risk level. General-purpose AI solutions may be evaluated:

- as a single complex use case – the Office should then take policy actions on the basis of the highest level of risk; or
- as separate use cases – the Office should then take policy action appropriate to each use case's respective risk, including appointing separate accountable use case owner and registering each in-scope use case on the internal register.[1]

Key Principles and Strategic Position

This policy has been developed in line with Australia's AI Ethics Principles[2]. These principles help us meet ethical standards, and ensure that our AI use is consistent with our strategic objectives, values and community expectations.

In summary the principles are as follows:

- **Human, societal and environmental wellbeing:** AI systems should benefit individuals, society and the environment.
- **Human-centred values:** AI systems should respect human rights, diversity, and the autonomy of individuals.
- **Fairness:** AI systems should be inclusive and accessible, and should not involve or result in unfair discrimination against individuals, communities or groups.
- **Privacy protection and security:** AI systems should respect and uphold privacy rights and data protection, and ensure the security of data.
- **Reliability and safety:** AI systems should reliably operate in accordance with their intended purpose.
- **Transparency and explainability:** 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.
- **Contestability:** 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.

---

[1] Policy for the responsible use of AI in government, p 7.

[2] Australia's AI Ethics Principles, Department of Industry, Science and Resources

- **Accountability:** People responsible for the different phases of the AI system lifecycle should be identifiable and accountable for the outcomes of the AI systems, and human oversight of AI systems should be enabled.

Decision makers within the Office should consider these Principles when making choices about how we will use AI. The principles should inform any decision making around AI, particularly in identifying and evaluating AI risks. The principles may also be used to interpret this policy.

The Policy aligns with the DTA Policy and reflects:

- The AI Ethics Principles
- The values articulated by staff and the executive during consultation
- The role, responsibilities and resourcing of the Office

Responsibility for AI in the Office

The agency Chief AI Officer is the Deputy Ombudsman who is responsible for leading AI adoption and associated change within the agency. This policy also provides that position with several other responsibilities.

Many roles within the Office have responsibilities under the policy, these are set out in the table below:

Role	Responsibilities
Chief AI Officer	Overall responsibility for leading the adoption and strategic change associated with AI implementation within the agency. Considers approval of medium and low risk AI use cases. Considers use case post implementation reviews Oversees remediation of any AI Incidents
All staff	Must comply with this Policy, the AI Acceptable Use Policy and complete any mandatory AI related training. Must report any AI-related incidents, hazards or unexpected behaviour.
Information and Technology Governance Committee (ITGC)	Provides strategic oversight of the development and implementation of information and technology policy, processes and systems across the Office. ITGC makes decisions on ICT and information management systems and infrastructure and approves major ICT releases.
Executive Committee	Considers the strategic and operational priorities of the Office, corporate governance, performance, resource allocation, and issues external and internal to the operations of the Office. Considers High Risk Use Cases for approval
Ombudsman	Accountable Authority and Executive Committee Chair
AI Accountable Official	The COO and CIO are the Office's designated AI Accountable Officials (AOs) under the Policy for responsible use of AI in government. They are also responsible for: • implementing this policy

	• Maintaining the Office AI Transparency Statement • Maintaining the Office AI Use Case Register • Ensuring the Office remains compliant with whole of government AI Policy. • Reporting to EC on AI use and risks • Notifying the DTA of any high-risk use of AI
Responsible Officer	Each AI use case approved by the Office has an accountable use case owner recorded in the AI Use Case Register They are responsible for: • applying the actions under the AI use case impact assessment • regularly monitoring and evaluating the use case • Taking action on AI Incidents reported to them. • ensuring their use case is registered with the Accountable Officials • applying the high-risk use case actions if the use case has an inherent high-risk rating • Conducting the Post Implementation Review on any use case they are responsible for.

Evaluation of Use Cases for the Office

The Office has adopted a risk-based process for evaluating any proposed new AI use cases before they are adopted. The process consists of five steps, these are:

1. Considering whether the AI Use Case is a prohibited use
2. Completing the Impact Assessment and Benefits assessment documents, including consulting with the key AI stakeholders
3. Deciding whether the AI Use Case should be approved.
4. Entering the approved use case on the AI Use Case Register and considering whether the Office AI transparency statement should be updated.
5. Conducting a post-implementation review

The Office uses the DTA's AI Impact Assessment tool to identify, assess, and manage AI impacts and risks against ethics principles. Officers completing the tool may find the accompanying guidance helpful.

The Office has decided to use the DTA tool to assess all AI use cases, excluding purely incidental uses of AI[3], and early-stage experimentation[4]. This includes use cases which do not fall within the scope of the DTA Policy. The level of detail and components of the assessment should be adjusted proportional to the risk, as explained below.

This evaluation process is used to obtain approval for the use of AI, however officers may also need to satisfy other policy requirements before a solution can be implemented:

- If the AI use case involves the deployment of a new technology, you will need to follow the Office System Lifecycle Policy, available here.

- If the use case involves the expenditure of public money you will need to follow an appropriate procurement policy. It is strongly recommended that officers refer to the Guidance on AI Procurement in Government when procuring AI products and services.

[3] Such as off-the-shelf software with AI features such as grammar checks and internet searches with AI functionality. The policy recognises that incidental usage of AI will grow over time and focuses on uses that require additional oversight and governance.
[4] Where that experimentation does not commit to proceeding with a use case or to any design decisions that would affect implementation later; risk harming anyone; introduce or exacerbate any privacy or security risks.

- If the use case involves any handling of personal information you may need to prepare a Privacy Threshold Assessment or a Privacy Impact Assessment. Contact the Legal team if you require advice.

Step 1: Prohibited uses

Some uses of AI are strictly prohibited by our Office. These uses are explained in detail below in the "prohibited use" section. These are:

- Discretionary Decisions
- Public correspondence, submissions and reports
- Unapproved technology and devices

Use cases only for a prohibited use must not be approved. A decision maker may approve a multi-purpose technology (such as Microsoft Co-Pilot). However, the approval must include a guardrail requiring staff not make prohibited use of the technology, and wherever possible preventing prohibited use.

Step 2: Impact Assessment Tool and Benefits assessment

The Responsible Officer for a new AI use case must complete the AI Impact Assessment tool, and the Benefits Assessment in the relevant templates. The officer should consider the supporting guidance to assist them to complete the tool.

This includes a detailed assessment of the benefits, risks, and planning process for the deployment of a new AI use case. The level of assessment required using the tool is proportional to the risk.

- Low-risk AI use cases need only be assessed against Sections 1 to 4 of the tool.
- Medium and high-risk use cases require a full assessment.

The Policy for the Responsible Use of AI in Government, does not apply to certain very low risk cases, however the impact assessment tool provides a convenient format for recording the details of a proposed AI implementation,

the identified risks and proposed treatments. The Impact Assessment Tool should be used to guide the assessment of all use cases within the Office[5]

For a very low risk use case, a Responsible Officer need only provide as much detail as is needed to properly evaluate the risks identified. This may be very limited in some cases, and Responsible Officers may record the information in a separate document, so long as they consider the same issues.

Expert Consultation

Expert engagement is critical to ensure that a technology being considered aligns with agency security and ICT policies, risk posture and data management policies. Part 1.10 of the AI Impact Assessment tool requires consultation with identified experts.

Key internal experts who should be consulted will depend on the nature of the use case. However at a minimum a Responsible Officer must consult the following individuals:

The AI Responsible Officers (CIO and COO)	All Use Cases
The Chief AI Officer (Deputy Ombudsman)	All Use Cases
The Chief Data Officer	Any use case involving the use of agency data
Directors – Legal	Any use case involving the collection or use of personal information
Directors -Policy and Strategic Insights	Any medium or high-risk use cases

This consultation process must be documented in the Impact Assessment.

Staff Consultation

Meaningful consultation with staff and unions is critical to building trust in AI adoption in the Office. It ensures employees have a voice in how AI is introduced, how to get the benefits, what problems can be solved with AI and

[5] With the exception of incidental AI use, and early stage experimentation as outlined above.

where it is likely to have a significant effect or material impact on them, including impacts related to gender, cultural identity, and First Nations peoples.

Consultation with staff is encouraged wherever the implementation of a use case may affect how employees do their work. It is required where an AI use case meets one of the criteria in cl 404 of the Office Enterprise Agreement 2024-27.

The APSC is expected to publish a circular regarding how and when to conduct consultation with employees regarding the implementation of AI in the Office. When this publication becomes available it should be considered during the consultation phase of the impact assessment process.

Consideration of risk controls

The Impact Assessment requires a Responsible Officer to document the inherent risk level of the proposed use case, indicate any planned treatments, and record the residual risk for medium and high use cases. When considering possible risk treatments an Officer should consider the following mitigation strategies in descending order of preference:

1. Changing internal processes to avoid identified risks arising.
2. Implementing technology controls to prevent identified risks
3. Implementing technology controls to monitor identified risks
4. Specific training to support Officers to avoid identified risks
5. Direction to staff not to use technology in certain ways.

These treatments are listed in descending order of strength. Wherever possible the Office should employ controls that prevent a risk from arising, over strategies that place responsibility with individual staff members.

Step 3: Consideration by the decision maker

The relevant decision maker must be provided with a decision minute attaching the Impact Assessment and Benefits Assessment setting out key information for their consideration. At a minimum either the Impact Assessment or the minute together must:

- Summarise the findings of the impact assessment

- Identify the Responsible Officer
- Identify any risks and set out any proposed risk treatments such as guardrails, limitations or controls required for the use case
- Identify what records generated by the use case need to be retained (if any) and how we will ensure compliance with record keeping requirements
- Identify any training required to ensure the use case can be delivered safely, and a plan for delivering that training
- Set out how information will be collected to monitoring effectiveness, bias and accuracy and to identify any AI Incidents
- Set out a plan to respond to any AI Incidents that arise out of the use case, and
- Set the date of the post implementation review.

In reviewing the Impact Assessment and considering whether to approve the proposed Use Case, the decision maker should have regard to:

- The Policy for the responsible use of AI in government
- The Office's Strategic Position on AI
- The AI Ethics Principles.

The appropriate decision maker is determined by the level of risk identified:

- **Low and medium risk** AI use cases may be approved by the Chief AI Officer. The Chief AI Officer may at their discretion refer the decision to the Information and Technology Governance Committee (ITGC) for consideration.
- **High risk** AI use cases must first be provided to the Chief AI Officer for review. A high-risk use case may only be approved after being considered by the Office Executive Committee. Additional controls exist for high-risk use cases, which are explained further below.

Step 4: AI Register and Transparency statement

AI Use Case Register

Once approved the Responsible Officer must ensure that the use case is recorded on the Office register of approved AI Use Cases, and that an up-to-date register is published internally on VOLT. A copy of the register must also be provided to the DTA every 6 months.

The template for the use case register is contained in **Appendix 2** of this Policy.

AI Transparency Statement

The Office is required to maintain an AI Transparency Statement that meets the requirements of the Standard for AI transparency statements. This is currently published on the Office's webpage here.

Before implementing an approved use case, the Responsible Officer should consider whether the Office transparency statement needs to be updated to reflect the new use.

Step 5: Post Implementation Review

The Responsible Officer for the use case must conduct a review between 3 to 12 months after the implementation, at the date set out in the decision memorandum. This review should include:

- Evaluating how the technology is being used.
- Whether evidence of bias, unexpected outputs or ethical issues has aligned with expectations.
- Confirming that all proposed training, proposed guardrails and controls have been implemented.
- Evaluating any complaints or incident reports collected during the operation of the system.
- Verifying the level of risk indicated in the original use case is appropriate.
- Recommending whether the Use case remains appropriate and setting out a plan for remediating any issues identified in the review.

This review should be set out in a memorandum and provided to the Chief AI Officer no more than 30 days after the commencement of the review.

High Risk Use Cases

If the Office's Executive Committee determines that an AI Use Case evaluated as "high risk" should be approved for use, this engages several additional responsibilities under the DTA Policy.

In addition to all other requirements the approval documentation must:

- Document that the use case has been reported to the Agency Accountable Official (the CIO and COO jointly)
- Document that the use case has been reported to the DTA through the Agency Accountable Officials.
- Determine a specific board or senior executive to govern the use case (recommended to be the Chief AI Officer)

- Determine a system to regularly review the use case at least every 12 months, including a report to the Executive Committee concerning whether the use case is operating as intended and whether risks are being effectively managed. The review must also consider the AI use case impact assessment and revisions to it, if required.

AI Incidents and Risk

All technologies sometimes fail to operate as expected. When AI does not operate appropriately it may create risks which could exist in any technology implementation, such as the risk of security breach, data loss or system outages. The CIO, CISO and ITSA are responsible for general ICT security and safety. These risks continue to be managed with our existing risk management framework.

AI may also create novel risks such as algorithmic bias or produce plausible incorrect output (sometimes called "hallucinations"). Planning for AI specific risks occurs during the approval process for an AI use case. A plan for identifying and responding to AI incidents should be documented in the approval memorandum for the use case, as set out below.

The appropriate level of planning depends on the nature of the use case:

- A use case which does not use any personal information, allows direct human oversight of each output, and doesn't directly interface with Office data (e.g. a text-to-voice tool, or a tool generating training materials) will only require minimal incident response planning. This may only require one or two paragraphs dealing with the matters outlined below.
- A use case which uses personal information, operates with some autonomy, or directly interfaces with Office data will require more detailed planning.

The exact appropriate measures depend on the circumstances, and should reflect the risks identified during the use case approval process. The incident response planning must cover the following matters:

- **Responsibility** – Who is responsible for identifying and actioning AI incidents in relation to this use case? This will usually be the Responsible Officer for the use case.
- **Detection** – How will we monitor for problems such as bias, "hallucinations" and data leakage in the use case? This could include human oversight of all outputs, sampling of outputs, period review, or automated monitoring.

- **Containment and Mitigation** – In the event an AI Incident occurs, what options does the responsible officer have to mitigate any damage, and how should they implement these quickly? This could include pausing the technology, disabling some components of a larger system, disconnecting an AI from sensitive data sources, or limiting the users who can access a system.
- **Notification** –AI Incident planning must determine who is responsible for notifying others of an incident, and who they should notify. Notifications must always include at least the following officers:

In these circumstances…	You must notify these officers…
All AI Incidents	Chief AI Officer, AI Accountable Officials, Chief Data Officer
All Privacy related incidents	Director – Legal

- **Remediation** –What options does the responsible officer have for remediating possible AI incidents?

While this planning process is separate, if the use case involves the deployment of a new technology, the Responsible Officer should attempt to align any AI Incident response planning with the ICT Incident Response Plan wherever possible.

Staff Identification of AI Concerns

Staff may identify a concern while using an AI system. This could be in relation to safety, security, reliability, bias, unexpected outputs or any other matter. Staff must report any AI Safety concerns to the Responsible Officer for the use case at first instance. The report must be made at the earliest opportunity after the concern is identifed.

If that person is not available, or cannot quickly be identified, they should report the concern to the AI Accountable Officials (CIO and COO).

Staff may choose to discuss a concern with a direct supervisor before reporting in the event of any uncertainty.

Staff should consider that all AI systems generate some plausible incorrect information. This would only be an AI safety concern where that information is exceeding the expected levels, or circumventing the controls put in place to ensure it does not result in harm.

The Responsible Officer or AI Accountable Officials on receiving a report from a staff member should consider:

- Whether it constitutes an AI Incident which requires engagement of the relevant response plan; if not
- Whether the reported issue should result in changes to the technology, use, safeguards or training provided.

Any staff identified concerns should also be considered in the post implementation review for the use case.

Complaints about the Office's use of AI

From time to time, the Office receives feedback from the public where they are not satisfied with the service they received. This may include complaints about the Office's use of AI or the raising of AI safety concerns.

Where members of the public make complaints about the Office's use of AI technology, staff should follow the Office's usual processes to record feedback about our services, as outlined in our *Policy – Feedback About our Service.*

Once recorded, any feedback regarding the use of AI or AI safety concerns should be provided to the Responsible Officer for the use case at first instance. If that person is not available, or cannot quickly be identified, the concern should be reported to the AI Accountable Officials (CIO and COO). They should consider:

- Whether the matter reported constitutes an AI Incident that requires engagement of the relevant response plan; if not
- Whether the reported issue should result in changes to the technology, use, safeguards or training provided.

Any public complaints should also be considered in the post implementation review for the use case.

Record Keeping

The Responsible Officer must establish a register of all AI Incidents, Staff concerns and public complaints in relation to an approved use case.

Documentation showing the evaluation of incidents and any action taken should be retained with this record. This must be stored with the approval minute for the use case on an authorised records system.

AI and Office Data

AI works most effectively when it is provided with high quality data inputs. AI implementation can also present a risk to the integrity of Office data. Because AI outputs are often opaque and non-deterministic, they can also present a risk to data quality if they are allowed to create or edit data.

The nature of the risks and the appropriate controls will vary depending on the nature of the use case. The AI Impact Assessment process outlined below requires an assessment of data quality risks posed by any particular use case being considered by the Office.

In all use cases that make use of, create, or edit Office data, the Responsible Officer for the use case is required to consult with the Chief Data Officer before seeking approval of the use case.

Prohibited use

Some uses of AI are strictly prohibited by our Office. AI systems must not be used for these purposes.

Discretionary decisions

Staff **must not** use AI to make decisions when exercising discretionary powers under any legislation. Discretionary decisions must always be made by staff with the relevant delegation, who consider the facts and circumstances of each matter individually.

Generally express legislative authority is required to automate decision making under an Act. No such authority exists in relation to any of our statutory powers.

Many actions in relation to a specific case or investigation involve the exercise of a specific power in our legislation and therefore cannot be actioned by AI. This includes:

- Deciding what action to take on any complaint, report, disclosure or notification
- Deciding on recommendations or factual findings in the preparation of a report.
- Deciding when and how to use our information gathering powers

Purely internal activities of an administrative nature, such as categorisation, allocation or triage, which do not change the outcome of the matter for

affected stakeholders are unlikely to be statutory decisions and do not fall within the scope of this prohibition.

AI may also support human decision makers to process information and identify relevant factors to consider in a decision. This is similar to a delegate's decision being supported by a memorandum prepared by staff who do not hold the delegation to exercise a power. For example:

- AI can summarise documents;
- AI can support officer research;
- AI can help decision makers identify relevant information for a decision;
- AI can help support analysis by identifying trends

A delegate using AI must ensure that they have not inadvertently delegated their decision-making authority:

- They must independently satisfy themselves of the circumstances relevant to the exercise of power, and must form any relevant state of mind before exercising their power.
- It is particularly important that a decision maker should have personally considered any arguments made by a party to the matter to ensure they have provided procedural fairness.
- AI should not be asked to recommend a specific decision. If a generative response AI suggests a decision unprompted, the user should document that they independently gave the matter consideration.
- A decision maker must ensure that they have kept records of how and why the power was exercised.

Public correspondence, submissions and reports

Staff **must not** use AI to directly prepare documents which will be sent outside the Office. This includes emails to parties, decisions, statements, social media posts, reports or submissions.

This does not prohibit the use of AI to locate relevant information, suggest a structure for correspondence, or to support analysis while preparing public communications. It specifically limits directly generating text or images for inclusion in communications.

For example a staff member must not use AI to directly generate a draft document, but could ask an AI tool to:

- "suggest a structure for this document"
- "suggest subheadings to break up the information in this report"
- "identify the key points I should include in the executive summary"

Staff may also ask AI tools to review and suggest improvements during the drafting process,

The use of AI in a public-facing document presents a reputational risk for the Office. It could create the perception that we have not given true independent consideration to a matter, which is integral to our functions. AI generated text frequently includes incorrect, but plausible information. These errors are often difficult for a proof-reader to identify.

The Office has decided we need to develop experience in the safe and effective use of AI before we consider employing it in generating material for use in any public-facing documents. This position may be adjusted in a future version of this policy as we become more sophisticated users of this technology.

Unapproved technology and devices

Staff **must not** use unapproved AI tools on Office devices, or for Office work. An up to date register of approved use cases will be available on VOLT in the event of any uncertainty.

The Office ICT team has taken steps to prevent access to many unapproved tools. However, it is the responsibility of staff not to access any unapproved AI tools, whether or not access is prevented. Staff should also be aware of the Direction from Home Affairs prohibiting the use of DeepSeek applications and web services.

Staff **must not** use any AI tool for Office work on a personal device. This includes accessing tools which are otherwise permitted under this policy. The Office network includes a broad spectrum of safeguards to ensure the safety of information that is provided to the Office. Using a personal device for Office work circumvents many of these controls and is strictly prohibited.

Misuse of AI will be handled under the Office's *ICT Acceptable Use Policy.*

Management of Records Created using AI

The Office manages records in accordance with the Archive Act 1983 and Information Management Standard for Australian Government. These obligations are technology neutral and do not distinguish between records created using AI and those created using other means.

The use of an AI technology might involve creating several new pieces of information, including:

- Prompts, instructions, parameters, and other inputs
- Output records and data; and
- Metadata such as records of the date, time and user for a set of prompts.

The Office follows the National Archives of Australia guidance on our obligations regarding records connected to AI –Information management for records created using Artificial Intelligence (AI) technologies.

The records we are required to retain are likely to be different for varying use cases. This should be considered as part of use case approval. Generally, AI is used to support and enable our work. In many cases AI will be used to provide a summary, support research, analyse voluminous data. These kinds of use are likely to be transitory and facilitative in nature and not represent the final version of a business record. In these circumstances we are not normally required to retain these records (although we may choose to do if they are useful to monitor the effectiveness of the use case).

For a general-purpose AI tool which may be used for many different purposes, some of these may generate records that require retention and others will not. In these circumstances it is acceptable to require users to distinguish between when they are creating facilitative, preliminary versions of a business record, and when they are creating a final version which must be retained. This is the approach taken with other general-purpose technologies used to create different records such as Microsoft Word, Microsoft Excel or Outlook.

If a Responsible Officer is unsure of the record retention requirements, they should consult the agency Records Manager during preparation of the Impact Assessment.

Training and Development

AI Training and developing our workforce is critical to help us realise much of the potential of AI in the Office. It is also an important part of managing many of the risks that may arise out of this emerging technology.

The Office will use a dedicated page within our knowledge management system (VOLT) to provide a central repository for internal AI guidance and training materials.

Training on responsible AI use is also mandatory for all staff under the *Policy for the responsible use of AI in government.* All staff must complete an AI fundamentals training module within 6 months of the commencement of this policy. The Chief AI Officer may also designate further training that must be completed by all staff members or cohorts in relation to AI in general, or particular use cases as required.

Policy Review

This policy will be reviewed annually to ensure it remains current and effective. An ad-hoc review may also be triggered by:

- a significant AI-related incident
- changes to relevant laws, regulations, or industry standards.

Appendix 1 – AI Use Case Register Template

Reference Number	Name of Use Case	Description	AI technology type	Accountable use case owner – name	Accountable use case owner - email	Life cycle status	Use of technical standard for government's use of AI

Domain	Usage pattern	Criteria the use case met to be in scope of the policy	Overall inherent risk rating	Overall residual risk rating	Date when AI impact assessment was last updated	Last date of review	Next date of review