AI USAGE POLICY

June 2026

Powering Skills Organisation Ltd

Policy # PSO-049	Last edited: June 2026	Policy owner: CEO
Effective date: June 2026	Next review date: June 2027	Approved by Board 30 June 2026

1. Purpose

This policy outlines the appropriate use of Artificial Intelligence (AI) systems by employees, contractors and stakeholders engaged with Powering Skills Organisation Ltd (PSO). The objective of the Policy is to ensure appropriate mechanisms are in place for the ethical and responsible use of AI in a work context by applicable stakeholders (see 2. Scope).

Specifically, it provides guidelines to safeguard PSO and stakeholder intellectual property (IP) and data. It also provides guidelines to ensure that outputs produced using AI systems are clear, actionable and validated before use. This policy will enable PSO to accelerate workforce intelligence and training product insights while meeting DEWR funding and reporting requirements and protecting the PSO brand and stakeholder trust in PSO.

Its usage is intended to enhance efficiency, insights, workforce intelligence, and operational productivity — including through the automation of routine tasks — while maintaining ethical, transparent and human-centred practices. AI outputs must be validated before use.

2. Scope

This policy applies to all PSO employees, contractors and PSO Board members using AI systems to create, automate, analyse, or communicate PSO content. This policy applies to external collaborators who use PSO data and/or outputs in creating outputs using AI systems. External collaborators must use AI systems that meet PSO's AI supplier conditions and must comply with PSO's external collaborator AI requirements as a condition of engagement.

This policy relates to any content created by, or that directly or indirectly incorporates outputs from, AI systems. This includes all uses of AI that produce or inform analysis, advice, decisions, models, reports, communications, or other outputs used by PSO in its professional, operational, or public-facing activities. It also covers correspondence, documents, social media content, and any other written, visual, or data-based media created fully or partially by AI systems, by any staff either for PSO's use or on behalf of, and/or for use by, any clients of PSO.

This policy does not relate to external committees (including but not limited to Technical Advisory Groups). Use of AI systems in the context of external committees will be governed via contracts, NDAs and transparency requirements.

3. Definitions

AI Accountable Executive: The role/person responsible for outcomes, compliance and reporting relating to AI usage. At PSO, this is the Chief Operating Officer.

AI-assisted content: Content where AI systems are used as a tool to support human creation. A human author retains primary control over the content structure, key messages, and final output; reviews, edits, and validates all AI-generated components; and takes professional accountability for accuracy and quality. AI-assisted content reflects substantial human authorship and expertise. Examples include using AI for drafting, editing, research summarisation, formatting, or data analysis.

AI-generated content: Content where AI systems autonomously create substantial portions of the output structure, key messages, or analysis. Human involvement focuses primarily on prompting, review, approval, and quality assurance rather than authorship. In accordance with Human-in-the-loop requirements, a nominated staff member must review, approve, and remain accountable for all AI-generated content before use. Examples include AI-written report sections, automated analysis summaries, or AI-created communications.

AI register: A central log of all AI tools and use cases within PSO. Content includes: owner; use cases; data classes; risk classification; pre-use test results; HITL arrangements; oversight arrangements; review dates; and, where applicable, the content and source documents used.

AI supplier conditions: The minimum requirements an AI system supplier must meet for their tool to be approved for PSO use, including but not limited to: (a) PSO data is not used to train the system; (b) PSO data is hosted and located entirely within Australia; (c) data provenance of PSO data is clearly documented; (d) retention and/or deletion windows for PSO data are clear; (e) the supplier provides transparency about AI model updates and significant changes; (f) the supplier meets applicable privacy and security standards; and (g) the supplier's contract terms are compatible with PSO's legal and funding obligations. Supplier conditions are reviewed annually or when a supplier makes material changes to their system.

AI system: A machine-based system that, for explicit or implicit objectives, infers from the input it receives how to generate output such as content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment. AI systems include but are not limited to: generative AI tools (such as Microsoft Copilot, Claude, ChatGPT), automated analysis tools, transcription tools, and machine learning systems, as documented in the Approved Tools Register.

API Key: Application Programming Interface key, used to identify and/or authenticate users. API keys must be given the same protective measures as any password or other security key.

Approved Tool: An AI tool that has passed PSO's governance process (privacy, security, supplier transparency, contract terms) and is listed in the Approved Tools Register.

Approved Tools Register: The central register maintained by PSO that documents all AI systems approved for use. It includes, for each tool: owner; use cases; data classes permitted; risk classification; pre-use test results; HITL arrangements; embedded AI features approved for use; and scheduled review dates.

Bias and fairness: The presence or risk of AI outputs that systematically disadvantage or misrepresent groups.

CEO: The Chief Executive Officer of PSO.

Content: Any content generated during the course of normal employment with PSO and includes but is not limited to analysis, reports, forecasts, communications or training product draft documents, spreadsheets, meeting transcripts and summaries, data visualisations and interactive tools, images, blogs and social media posts, educational materials and promotional materials.

Data minimisation: Collecting, processing, and storing only the minimum amount of data necessary for a specific purpose.

Data provenance: The record of where data comes from, how it has been collected, and what transformations or validations have been applied.

De-identified data: Data that has had all personal identifiers removed so that individuals cannot reasonably be re-identified.

DEWR: The Federal Government Department of Employment and Workplace Relations.

External collaborators: Individuals or organisations who are not PSO employees, contractors, or Board members, and who access, use, or contribute to PSO data, outputs, or work products, or create deliverables for PSO or PSO clients using PSO resources or information. This includes but is not limited to consultants, subject matter experts, partner organisations, and subcontractors engaged for specific projects or deliverables. It excludes members of external committees (e.g. Technical Advisory Groups) whose AI use is governed by separate contractual arrangements.

Funding Program: Any Federal, State or Territory government program that provides funding for/to PSO.

High-risk use: An AI use case where outputs (Content) may materially influence funding, policy, training products, workforce strategy, or decisions about people. See also: Risk classification.

Human-in-the-loop (HITL): A system or process in which a human actively participates in the operations, supervision, or decision-making of an automated system. In the context of this policy, HITL means that humans are involved at some point in the AI workflow to ensure accuracy, safety, accountability or ethical decision making. PSO will nominate a staff member with the responsibility to review, approve, and remain accountable for AI-influenced outputs.

ICT: Information and Communications Technology.

Kill-switch: The immediate suspension of an AI system's use by the AI system owner. Kill-switch conditions include but are not limited to: detection of a data breach involving the AI system; the system producing systematically biased or harmful outputs; a model update causing material accuracy degradation; or direction from the AI Accountable Executive or CEO.

Notifiable Data Breach (NDB) Scheme: The scheme established under the Privacy Act 1988 (Cth) that requires organisations to notify affected individuals and the Office of the Australian Information Commissioner of eligible data breaches.

Personal information: Any information that can be used to identify a person, either alone or when combined with other information, as specified in the Privacy Act 1988 (Cth), and includes but is not limited to name, age, address, email, phone, financial, and other information.

Policy: This AI Usage Policy.

Policy Officer: Whoever is charged with the role of Policy Officer at PSO, as may change from time to time.

Prompt hygiene: Practices for crafting AI system inputs that minimise security and privacy risks, reduce bias, and achieve accurate outputs. This includes: not including sensitive personal identifiers unnecessarily; being specific about required outputs and their intended use; applying data minimisation before submitting any data to an AI system; and validating AI responses before use or publication.

Proprietary data/Confidential data: Information belonging to PSO, DEWR, or partners that is not public, including contracts, funding details, unpublished training products, or sensitive workforce insights. For the purposes of this policy, 'confidential' information includes all data classified as Confidential or Highly Confidential under PSO's data classification framework (see PSO IT Usage Policy PSO-044). This includes, but is not limited to, personally identifying information (PII) such as names combined with contact details, financial information, health or WHS records, government datasets, and API keys.

Risk classification: The categorisation of AI use cases by their potential to cause harm. For the purposes of this policy: Low risk — internal-only use, no sensitive data, outputs reviewed before use. Medium risk — external or public-facing use, or use involving internal data, requiring approver sign-off and HITL review. High risk — any use case where outputs may materially influence funding, policy, training products, workforce strategy, or decisions about people, requiring detailed risk assessment, named HITL nominee, and documentation in the AI Register. Data classification categories are defined in the PSO IT Usage Policy (PSO-044): Public, Internal, Confidential, and Highly Confidential.

Staff: Any individual working for PSO, including but not limited to full-time, part-time, permanent or casual, contractor, employee or volunteer.

VET: Vocational Education and Training.

4. Guiding Principles

PSO commits to accountability, transparency, meaningful human oversight (Human-in-the-loop), risk management, and fairness. Where data relates specifically to Aboriginal and Torres Strait Islander peoples, PSO will ensure it is governed with appropriate consultation and respect for community expectations.

Risk Assessment and Compliance

PSO's appetite for risk in terms of AI usage, and ICT/cyber matters generally, is low. Compliance with applicable legislative instruments and the regulatory frameworks governing PSO comes first. This includes the Privacy Act 1988 and Australian Privacy Principles, the DEWR Third-Party AI Assessment Framework, Funding Program rules, and the VET Quality Framework. See also Related Documents in this policy.

Innovation involving the use of AI systems must occur strictly within the controls of this policy. Prior to the use of an AI system, a risk check must be completed and evidence retained, and the activity captured in the AI Register.

Accountability and Attribution

AI system use cases and approved tools must be logged in the AI Register, which is owned by the Chief Operating Officer and maintained, where practicable, through an automated log in the PSO SharePoint AI Register location. Routine day-to-day use of approved tools (such as rephrasing text, grammar checking, drafting internal documents, or generating internal visuals) does not require individual interaction logging unless the use contributes to: (a) conditional or high-risk uses; (b) new use cases not previously risk-assessed; (c) content for external publication or Board or Committee consideration; or (d) any use where the content and relevant source document should be captured for auditability. The AI Register entry must identify the relevant system owner, approver and nominated HITL staff member responsible for oversight. Accountability for decisions and outputs remains with PSO personnel and cannot be delegated to the AI system.

Transparency

PSO must declare the use of AI systems in public outputs or training materials produced by or on behalf of PSO. PSO will publish a public statement on its website describing how it uses AI systems. AI system use will be auditable and FOI-ready through the records maintained in the AI Register. See Section 9 for detailed disclosure requirements.

Integrity

The use of AI systems must not undermine the authenticity or integrity of PSO's analysis or products. Authenticity means that PSO's outputs accurately represent the organisation's evidence-based analysis, the underlying data and methodologies used, and any limitations of that analysis — including where AI systems have contributed to the output. AI use must not be used to manufacture, misrepresent, or inflate the evidential basis of PSO's work.

All staff are required to verify AI-generated and AI-assisted content before it is used in reports, decisions, or communications. Verification means reviewing the output for accuracy, relevance, potential bias, and compliance with PSO's data and ethical obligations — and not simply accepting the AI system's output as authoritative.

Safety and Equity

Human oversight activities must ensure that AI system outputs are not biased against equity groups, including but not limited to women and First Nations. First Nations content and insights produced by or with involvement of AI systems must be handled in a culturally appropriate manner.

5. Acceptable and Prohibited Uses

Use of AI systems is divided into Acceptable, Conditional and Prohibited uses. Only Approved Tools can be used, regardless of use category. Authorised users are only permitted to access AI systems to which they have been granted access rights. Access rights are granted on the basis of business needs and position description. See the PSO IT Usage Policy (PSO-044) for further information on access to systems.

This table is not exhaustive. New use cases may be assessed using PSO's risk classification framework and, following approval via the change request pathway in Section 6, added to the appropriate category in the AI Register. All uses, regardless of category, require HITL oversight as defined in Section 3. The categories below determine the level of additional approval and disclosure required.

Acceptable uses require human review of Content (outputs). Generally, this relates to any internal-only use or Content.

Conditional uses require prior written approval from the Chief Operating Officer (or delegate) and disclosure of AI use in the final output. Generally, this relates to any external or public-facing use or Content. Approval must be documented and retained in the AI Register.

Prohibited uses are not permitted without exception under normal circumstances. In exceptional circumstances where a business-critical need requires exploration of a use case that would otherwise be prohibited, written approval must be obtained from the CEO and documented in the AI Register with a dated risk assessment. This pathway is not available for uses that would breach Australian law or PSO's funding obligations.

Acceptable	Conditional	Prohibited
Drafting internal-use only content	Public reports	Use of AI systems not on the Approved Tools list
Analysis and research assistance	Board reports (with strong HITL requirements)	Uploading Confidential/Highly Confidential/DEWR data to unapproved tools without completing a risk assessment
Summarising consultations	Media lines	Uploading materials with a government security marking (e.g. SEC:OFFICIAL) without AI Accountable Executive approval
Scenario ideation	Stakeholder communications	Sole reliance on AI for workforce forecasts, labour market data, or board/policy decisions without independent expert validation
Query and information assistance	Training package content development	Automating external decisions about people, funding or policy
Formatting and grammar checking	Consultation reports for external use	Generating compliance advice without expert review
Meeting note transcription and summarisation	Recording internal meetings and using AI tools to generate draft minutes or action summaries, where recording is announced at the commencement of the meeting, attendees agree, and the output is human reviewed	Creating training assessment materials without validation
Data visualisation creation	Recording Board, Committee, external stakeholder or otherwise sensitive meetings and using AI tools to generate draft minutes or summaries, subject to prior approval, announced recording at the commencement of the	Sharing API keys or system credentials via AI systems

	meeting, attendee agreement, human review, and deletion of recordings and AI tool artefacts once final minutes are confirmed in accordance with PSO requirements	Using AI to generate content that misrepresents PSO's position or creates false attribution
		Circumventing security controls or data protections Non work-related use of AI systems that is inconsistent with the PSO Code of Conduct, PSO IT Usage Policy, approved access conditions, or PSO's privacy, security and records management requirements

6. Governance, Procurement and Approvals

Governance

The AI Accountable Executive is the Chief Operating Officer. Management will maintain an internal AI Council to support implementation of this policy, review proposed tools and use cases, monitor emerging risks, and prepare matters for escalation through PSO's governance pathway. The AI Council operates under the authority of the AI Accountable Executive and does not replace the approval roles of the Audit and Risk Committee, CEO or Board.
The Audit and Risk Committee is to function as an ongoing governance forum for AI usage. Management through the AI Council will provide periodic reporting to the Audit and Risk Committee and the Board on approved tools, material use cases, incidents, emerging risks, monitoring outcomes and proposed policy updates. Where the requester or business owner of a proposed tool or use case is ordinarily part of the approval chain, the matter must be escalated to an independent approver or governance body to manage conflicts of interest.

Procurement and Approvals

AI systems cannot be used prior to approval. A list of approved AI systems is maintained in the Approved Tools Register.

Change requests relating to AI usage are to follow this approval pathway:

Initiated by relevant Team Leader → Agreed by COO → Reviewed by Audit Risk Committee → Approved by PSO Board

AI-related change requests include but are not limited to: requests for additions to the Approved Tools Register; new or changed Acceptable, Conditional or Prohibited Uses; changes to Supplier Requirements; changes to the content or format of the AI Register; changes to the content or format of the Approved Tools Register; changes to the AI Usage Policy. The AI Register is owned by the Chief Operating Officer and stored in the designated PSO SharePoint AI Register location. Where practicable, the register is maintained through an automated log to support auditability and record integrity.

Supplier Requirements

Vendors of AI systems must meet the following minimum requirements to be added to the Approved Tools Register. Additional requirements and/or exceptions can be imposed by the Audit Risk Committee or PSO Board during the review and approval process:

1. PSO data must not be used to train the system.
2. PSO data must always be hosted and located entirely within Australia.
3. Data provenance of PSO data must be clearly known.
4. Retention and/or deletion windows of PSO data must be clear.
5. The supplier must provide transparency about AI model updates and significant system changes.
6. The supplier must meet applicable privacy and data security standards.

API-Based AI Services

AI services accessed via Application Programming Interface (API) within PSO-developed or PSO-managed applications may be approved where all of the following conditions are met:

- The API provider's terms confirm PSO data is not used for model training.
- The API provider specifies data retention and/or deletion windows (e.g. zero retention or maximum 30-day retention).
- PSO data is stored in PSO-controlled environments (Azure/SharePoint) hosted in Australia.
- Data is only transmitted to the API for processing and is not permanently stored by the vendor.
- The PSO-developed application implements appropriate security controls and access restrictions.
- API keys are managed securely and treated with the same protections as passwords.

Use of API-based AI services must be approved by the AI Accountable Executive, documented in the AI Register with an assigned HITL nominee, and assessed under the relevant Supplier Requirements before activation.

7. Pre-use Testing and Ongoing Monitoring

Records of pre-use testing, records of approvals, and evidence of ongoing monitoring are to be stored in the AI Register.

Pre-use Testing

Before use, each AI system will pass documented acceptance tests. Acceptance tests will vary from system to system and use to use. The Audit Risk Committee, in consultation with relevant subject matter experts, will determine and document the appropriate acceptance test criteria during the approval process.

Acceptance criteria may include accuracy bounds, bias checks, accessibility reviews and latency/reliability threshold limits. Guidance by tool type:

- Text generation tools: Accuracy testing against known outputs, bias detection using representative test scenarios, verification of no-training commitment.
- Data analysis tools: Accuracy testing against manually calculated outputs, consistency checks, edge case testing.
- Transcription tools: Accuracy testing against known transcripts, handling of PSO-specific industry terminology.

Ongoing Monitoring

The AI system or use owner, as documented in the Approved Tools Register and/or AI Register, will be accountable for monitoring the performance of the AI system. Monitoring arrangements, including cadences and/or triggers (e.g. model updates), rollback triggers, kill-switch triggers, incidents, exceptions, and material changes to use cases are to be documented for each AI system by the AI system owner.
Monitoring must occur at a minimum quarterly, or more frequently based on risk level and usage volume. High-risk AI systems must be monitored at least monthly. Monitoring findings, material incidents, approved exceptions, and any proposed control changes must be recorded in the AI Register and reported through the governance pathway at a frequency determined by the AI Accountable Executive, with material matters escalated promptly to the Audit and Risk Committee and, where required, the Board.
The AI system or use owner, as documented in the Approved Tools Register and/or AI Register, will be accountable for monitoring the performance of the AI system. Monitoring arrangements, including cadences and/or triggers.

8. Data Governance, Privacy and Security

PSO will apply data minimisation, require provenance documentation for datasets, and contract no-training-on-our-data, with Notifiable Data Breach (NDB) Scheme-aligned breach response processes.

Forbidden Data Classes

Prohibited — Third-Party AI Web Interfaces:

Users are prohibited from uploading the following data classes to third-party AI web interfaces (e.g. ChatGPT.com, Claude.ai, Copilot web interface when used outside PSO's Microsoft 365 tenancy):

- Personally Identifying Information (PII)
- Confidential consultation records
- Unpublished training product drafts
- Embargoed reports
- Contracts
- DEWR-provided data
- Materials carrying a government security marking (e.g. SEC:OFFICIAL or equivalent)

Restricted — PSO-Controlled Applications Using Approved APIs:

The above data classes may only be processed through PSO-developed or PSO-managed applications using approved AI APIs (as defined in Section 6) where all of the following apply:

- A risk assessment has been completed and approved by the AI Accountable Executive.
- The API service meets all Supplier Requirements including any API-specific criteria.
- The use is documented in the AI Register with an assigned HITL nominee.
- Data minimisation and de-identification techniques have been applied where practicable.

Two-Tier Data Upload Risk Assessment

All staff must apply the following risk assessment before uploading any data to an AI system:

Tier 1 — Quick Pre-Use Check (mandatory for all data uploads to any AI system):

Before uploading any data, staff must confirm: Is this an Approved Tool? Does the data classification of this information permit use with this tool? Is this the minimum data needed for the task?

Tier 2 — Detailed Risk Assessment (triggered where Tier 1 reveals any uncertainty, or where data is Confidential or Highly Confidential):

A documented risk assessment approved by the AI Accountable Executive before proceeding. Uploading organisational information, personal information, or sensitive information (including Confidential, Highly Confidential, proprietary, or DEWR data) to any AI system — approved or unapproved — without completing the appropriate Tier check remains prohibited. Staff must apply data minimisation, de-identification, redaction and classification checks before any upload, and must not include personal or sensitive information unless the approved use case, tool controls and risk assessment explicitly permit it.

Data Minimisation and Anonymisation

Staff must apply data minimisation techniques when working with AI systems, including applying anonymisation and redaction to datasets before submission. PSO will provide staff with approved anonymisation and redaction templates and guidance on their use. Training on these templates will be included in AI onboarding and refresher training. Where no suitable template exists for a dataset, the AI Accountable Executive must be consulted before proceeding.

Data Provenance

Data provenance records must be kept in the AI Register, including any limitations of the dataset(s).

Data Security

AI systems integrated within PSO's Microsoft 365 environment are to inherit SharePoint/tenant permissions, and access is not to be broadened. Retention and deletion schedules, as documented in the AI Register, are to be observed.

For standalone applications used by PSO that incorporate AI features (such as Borealis or Converlens), the AI Accountable Executive and IT Service Desk (Novigi) must assess and document: (a) what data those AI features can access; (b) how data is processed and stored; and (c) whether use of those AI features is approved under this policy. Staff must not activate or use AI features in standalone applications without confirmation that they have been assessed and approved.

In the event of an AI system-related Notifiable Data Breach event, the response process documented in the PSO Information Management Plan (PSO-041) is to be followed.

First Nations Data

First Nations content/insights produced by or with the involvement of AI systems must be handled in a culturally appropriate manner. Queries regarding culturally appropriate handling of applicable content/insights should be directed to the CEO or a nominated First Nations cultural authority.

9. Disclosure and Labelling

In keeping with the Guiding Principle of Transparency, PSO will publish a public statement on its website describing how PSO uses AI systems. This statement will serve as the primary disclosure mechanism for external stakeholders and will be reviewed and updated annually or when AI usage materially changes.

AI system use will be auditable and FOI-ready through the records maintained in the AI Register.

Individual documents are not required to carry AI disclaimer text as a default. Where AI use is particularly material to a specific output and the document owner — in consultation with the AI Accountable Executive — considers a contextual note appropriate, they may include one. Disclosure wording must be clear, accurate, and proportionate to AI's actual role in the content's creation.

Stakeholders wishing to understand how AI was used in a specific PSO output should be directed to the PSO AI public statement on the PSO website or to the AI Accountable Executive.

10. Contestability and Stakeholder Engagement

In keeping with DEWR's Integrity Framework, which enshrines transparency as the default position for all actions, PSO will ensure that outputs produced by or with the involvement of AI systems can be challenged or queried by stakeholders.

The AI Accountable Executive (Chief Operation Officer) is the single point of contact for all external enquiries relating to PSO's AI use, regardless of the specific document or product. This ensures high-level visibility of the scale and type of enquiries received, and that responses are developed in considered consultation with appropriate staff.

The document/product owner will remain accountable for ensuring pre-publication reviews for bias and accessibility occur (as documented in the AI Register).

AI-related issues are to form a part of routine stakeholder engagement with industry, unions, RTOs and communities.

11. Training and Capability

All staff must complete mandatory AI awareness training as part of onboarding before using any AI system for PSO work. This training will cover safe use, privacy obligations, disclosure requirements, prompt hygiene, output verification, and use of PSO's anonymisation and redaction templates.

Refresher training will be delivered at a minimum annually, or sooner when: (a) AI systems in use materially change; (b) significant incidents occur; or (c) relevant legislative or regulatory guidance is updated.

AI-system approvers and staff responsible for monitoring will receive further training on pre-use testing, risk assessment, and ongoing monitoring responsibilities (see Section 7).

Executives and PSO Board members will receive briefings at least annually and as significant AI risks, incidents, or regulatory changes emerge. Briefings will cover: current and planned AI tool usage at PSO; material AI risks and incidents; changes to the regulatory environment; and updates on PSO's compliance with this policy. The AI Accountable Executive is responsible for preparing and delivering these briefings.

AI-related training will be tracked and recorded in the PSO Learning Management System (LMS).

New users must acknowledge this policy upon onboarding in the Human Resources Information System (HRIS).

12. Review Cycle and Continuous Improvement

This Policy is to be reviewed annually. Data and records from pre-use testing, ongoing monitoring, incidents, exceptions, training completion, and the AI Register are to be considered during the annual policy review. Proposed updates must be assessed through the governance pathway in Section 6 and approved by the relevant approval authority.

Additional reviews may be triggered by: AI system changes, including model or vendor changes, significant incidents or near misses, emerging legal, regulatory or funding guidance, material changes in PSO’s risk appetite or operating environment, repeated exceptions or non-compliance, or recommendations from the AI Council, Audit and Risk Committee or Board. The policy owner, in consultation with the AI Accountable Executive, is responsible for initiating annual and ad-hoc reviews.

13. Roles and Responsibilities

Position	Responsibility
CEO	Overall governance of AI Usage policy and compliance. Approves AI-related change requests escalated by the PSO Board. Provides CEO-level approval for any exceptional use of an otherwise-prohibited AI function.
Chief Operating Officer (AI Accountable Executive)	Policy implementation and enforcement. Owns the AI Register, including oversight of its maintenance in the designated PSO SharePoint AI Register location and, where practicable, through an automated log. Single point of contact for all external AI enquiries. Prepares and delivers AI governance briefings to executives and the PSO Board. Approves conditional use requests and risk assessments for restricted data use. Reviews and approves API-based AI service proposals.
Audit and Risk Committee	Ongoing AI governance forum. Reviews and approves AI-related change requests. Determines acceptance test criteria in consultation with subject matter experts. Oversees policy compliance and monitors AI-related risks.
PSO Board	Approves AI-related change requests with an identified risk tolerance and outside of the scope of the identified and agreed approval pathway submitted via the governance pathway. Receives annual AI governance briefing from the AI Accountable Executive.
Policy Officer	Maintains this policy. Coordinates annual and ad-hoc reviews. Monitors legislative and regulatory changes affecting the policy and alerts the AI Accountable Executive.
IT Services Desk (Novigi)	System administration, security, monitoring, and incident escalation. Assesses AI features in standalone applications. Manages Approved Tools Register in conjunction with the AI Accountable Executive.
Team Leaders	Initiate AI-related change requests. Support staff compliance with this policy. Conduct 'is this in Copilot/approved tools?' conversations with team members. Escalate concerns to the General Manager.
AI System Owners	Accountable for monitoring the performance of individual AI systems. Document monitoring arrangements in the AI Register. Initiate rollback or kill-switch where required. Conduct or commission pre-use testing.
Human-in-the-Loop (HITL) Nominees	Nominated staff responsible for reviewing, approving and remaining accountable for AI-influenced outputs in their area, before those outputs are used, published, or submitted externally.
Document/Product Owners	Responsible for disclosure decisions on individual outputs. Accountable for pre-publication reviews for bias and accessibility. Direct stakeholder AI enquiries to the AI Accountable Executive.

Authorised Users	Comply with this policy. Apply the two-tier data upload check before submitting data to any AI system. Report incidents and suspected policy violations. Safeguard data and AI system credentials.

14. Compliance

Suspected or actual non-compliance with this policy must be identified, documented and escalated through PSO's management processes. The AI Accountable Executive is responsible for ensuring that material breaches, incidents, repeated exceptions, or control failures are investigated and, where appropriate, reported to the Audit and Risk Committee and the Board. Responses may include additional supervision, training, suspension of tool access, withdrawal of approvals, remedial actions, disciplinary action, or other employment or contractor consequences in accordance with PSO policies and applicable law. Use of an AI system may also be suspended immediately where a kill-switch trigger, data breach risk, unlawful use, or material control failure is identified.

15. Related Documents

- Privacy Act 1988
- Australian Privacy Principles
- DEWR Third-Party AI Assessment Framework
- Notifiable Data Breaches Scheme
- VET Quality Framework
- Invasion of Privacy Act 1971 (Qld)
- Telecommunications (Interception and Access) Act 1979 (Cth)
- Workplace surveillance legislation
- PSO Code of Conduct
- PSO Privacy Policy
- PSO IT Usage Policy (PSO-044)
- PSO Information Management Plan (PSO-041)
- PSO Information Security Policy
- PSO Business Continuity Plan
- PSO Approved Tools Register
- PSO AI Register

16. Document Control Table

Version	Description	Date
0.1	First draft	September 2025
1.0	Approved version — October 2025	October 2025
3.0	Revised following PSO AI Council review, Board and Audit and Risk Committee feedback, and staff consultation. Includes governance, monitoring, review cycle, meeting recording, accountability and enforcement clarifications.	March 2026
3.1	Minor revision following April 2026 ARC ahead of Board endorsement	June 2026