Acceptable Usage of Generative Artificial Intelligence (AI)

Status:	Final
Issue Date:	14/05/2026
Last Review Date:	14/05/2026
Last Approval Date:	14/05/2026
Last Publish Date:	14/05/2026
Review Frequency:	6 monthly (or as required)
Document Version:	1.1
Responsible Division:	ICT
Contact:	Chief Information Officer
Type:	Policy
Document Classification:	Internal

DOCUMENT SIGN-OFF

Job title	Role	Date
Chief Information Officer	Approver	14/05/2026
Head of Security	Document Owner	14/05/2026

1 PURPOSE

This policy governs the use of Generative Artificial Intelligence (GenAI) tools and platforms to ensure their use supports LRS's objectives while managing potential risks to data security, privacy, intellectual property, operational reliability, and compliance with LRS's concession obligations. This policy has been defined by the LRS Artificial Intelligence Steering Committee (AISC).

This policy applies to all LRS employees, directors, contractors, and third-party partners when using any AI tools or platforms to perform any work on behalf of LRS - whether on LRS premises, remotely, or on personal devices.

This policy applies to:

- Standalone GenAI tools approved by the AISC (e.g. Microsoft Copilot for Enterprise, approved LLM platforms)
- Publicly available, free-tier, or consumer AI tools used for permitted generic tasks
- AI features embedded in enterprise software platforms (e.g. Adobe Acrobat, AI mode in Google Chrome)

This policy is subordinate to the LRS AI Governance Framework and must be read alongside it. Defined terms used in this policy carry the meanings set out in the Governance Framework Glossary (Section 2 of that document).

2 CORE PRINCIPLES

This policy is grounded in the following principles, which are consistent with the LRS AI Governance Framework core principles:

Human-Centred Accountability - Individuals remain responsible for all advice, decisions, and output generated or assisted by AI. AI tools are aids, not decision-makers. The fact that AI generated or contributed to an output does not reduce individual accountability for that output.

Transparency - All AI-influenced or AI-generated content must be clearly identified as such when it is used in official communications, records, reports, or formal advice. Where AI has materially contributed to a piece of work, this must be disclosed.

Fairness - Outputs must be checked for bias and inaccuracy before use. Users must not rely on AI outputs that could produce outcomes that negatively impact individuals, LRS, or the public.

Safety - All use of GenAI tools and platforms must protect the safety and security of LRS systems, data, staff, customers, and the public. This includes the security of data submitted to AI tools as prompts or inputs.

Accuracy - Users must not treat AI outputs as authoritative without verification. For tasks involving legal rights, property records, official figures, or regulatory obligations, outputs must be verified against primary sources by someone with relevant competence.

HUMAN IN THE LOOP - ALWAYS	NO AUTOMATED DECISION-MAKING
Every AI output must be reviewed by a qualified human before any action is taken. AI is a tool to assist human judgement, never to replace it.	AI must not make or execute final decisions on any matter affecting individuals, property records, or LRS obligations. Humans must make every decision.

3 AI TOOL CATEGORIES

All AI tools fall into one of three categories for the purposes of this policy:

Category	Description	Data Permitted
Category A - AISC-Approved Enterprise AI	AI tools or platforms formally approved by the AISC for processing LRS data, with contractual data protection obligations in place.	Sensitive or Internal data (as defined). Only for approved use cases.
Category B - Publicly Available AI Tools	Consumer-grade or free-tier AI tools (e.g. free ChatGPT, Google Gemini consumer, Claude.ai). Not approved for sensitive or internal data. Permitted only for approved generic tasks.	De-identified or publicly available information only. Never Sensitive or Internal data.
Category C - Shadow AI	Any AI tool or service used without AISC knowledge or approval, including personal AI subscriptions used for work purposes.	Prohibited. No LRS data of any kind.

Note on enterprise-licensed versions: An enterprise licence for a tool (e.g. ChatGPT Enterprise, Microsoft Copilot for M365) does not automatically make it a Category A tool. It must be formally submitted for AISC approval. An enterprise licence may be a prerequisite for approval but is not sufficient on its own.

4 ACCEPTABLE USAGE RULES

The following three core rules govern all use of AI tools by LRS employees, contractors, and partners. These are rules, not principles - compliance is mandatory.

4.1 RULE 1: USE ONLY APPROVED OR PERMITTED TOOLS

- Use only Category A (AISC-Approved) tools for any task involving Sensitive or Internal LRS data.
- Category B (Publicly Available) tools may only be used for the permitted generic tasks listed in Section 6.
- Category C (Shadow AI) is prohibited for all LRS work purposes.
- If you are unsure which category a tool falls into, contact the IT Service Desk or CIO before use.

4.2 RULE 2: APPLY PROPORTIONATE HUMAN OVERSIGHT

All AI-generated output must be reviewed by a human with relevant competence before being used in any official capacity. The level of review required is proportionate to the risk tier of the use case:

- Risk Tier 1 (Low): User self-review is sufficient.
- Risk Tier 2 (Moderate): Mandatory review by the user; factchecking against a verifiable source required before use in any official document or communication.
- Risk Tier 3 (High): Expert verification required. The reviewer must be competent to assess the accuracy and completeness of the output independently of the AI tool.
- Risk Tier 4 (Critical): Committee-level or senior management sign-off required. AI output must not be used without documented approval.

4.3 RULE 3: OBSERVE ALL PROHIBITED ACTIVITIES

The following activities are expressly prohibited under this policy:

Shadow AI - Using any unvetted, unapproved, or personal AI account or tool to process, analyse, or draft any LRS-related work. This includes exporting LRS data to a personal ChatGPT, Claude, Gemini, or similar account, processing it using the AI tool, and re-importing the result into LRS systems. This practice is prohibited regardless of whether the output 'looks right' or the data is considered non-sensitive by the individual. It bypasses all security, audit, and compliance controls.

Exposing Sensitive or Internal Data to Unapproved Tools - Uploading, entering, copying, pasting, or summarising any Sensitive or Internal data (as defined in Section 7) into any Category B or Category C AI tool. This includes entering data via prompts, image uploads, or file attachments.

Unsupervised Automated Decision Making - Using AI to make or execute final decisions on matters that affect individuals' rights, entitlements, or property records, or that have significant operational or financial consequences for LRS, without documented human sign-off. A decision is 'significant' if it: (a) affects a legal

right or entitlement; (b) has a financial value above [THRESHOLD - to be set by AISC]; or (c) is irreversible without material cost or effort.

Jailbreaking - Attempting to bypass, circumvent, or manipulate the built-in safety filters, ethical guardrails, content restrictions, or security controls of any AI tool. This includes prompt engineering specifically designed to elicit prohibited outputs.

Training External Models on LRS Data - Submitting LRS data - including via prompts - to any AI service that uses submitted data to train or improve its underlying model, unless explicitly approved by the AISC with contractual data rights protections in place.

5 APPROVED USE CASES FOR CATEGORY A (AISC-APPROVED) AI TOOLS

There are currently no Category A (AISC-Approved) GenAI tools authorised for use with Sensitive or Internal LRS data. This section will be updated as tools are approved through the AISC process.

Business units wishing to propose a tool for Category A approval must submit an AI Use Case Request to the AISC (see the AI Governance Framework, Section 10, for the request template and process). The AISC has committed to providing a decision within 30 business days for all formal submissions.

In the interim, teams requiring AI assistance with Sensitive or Internal data must consult with the CIO to assess options. The CIO may authorise interim supervised use of specific tools under defined conditions pending AISC review.

6 PERMITTED USE CASES FOR CATEGORY B (PUBLICLY AVAILABLE) AI TOOLS

Category B (Publicly Available) AI tools may be used for the following task types, subject to the guardrails in Section 8. These tools must never be used to process Sensitive or Internal data.

6.1 CONTENT CREATION AND REFINEMENT

- Drafting or improving tone and grammar for external correspondence, where no LRS-specific content is included
- Grammatical assistance and spell-checking on publicly available or fully de-identified text
- Generating template structures or example formats for internal documents, where no LRS data is included

6.2 TECHNICAL ASSISTANCE (DE-IDENTIFIED ONLY)

- Generating generic code snippets, SQL query patterns, or Excel formulas for internal technical use, subject to expert human review before execution
- Summarising publicly available industry reports, white papers, or news articles
- Troubleshooting common operating system or software errors using publicly available information

6.3 RESEARCH AND IDEATION

- Researching external, publicly available data and information
- Translating publicly available documents or standard phrases

Brainstorming ideas or options where no LRS-specific or confidential context is shared

Using AI for analysis or decision support at a general, non-LRS-specific level, where outputs are treated as preliminary and subject to human expert review

6.4 LEARNING AND DEVELOPMENT

- Summarising external educational materials and articles
- Role-playing professional conversations or client negotiations using entirely hypothetical, non-specific scenarios containing no real client details, case references, or LRS-identifiable information.

7 DATA CLASSIFICATION: SENSITIVE OR INTERNAL DATA

Sensitive or Internal Data includes, but is not limited to, the following categories. When in doubt, treat data as sensitive.

Data Category	Examples
Land Titles and Property Data	Title records, dealings, plans, caveats, mortgages, easements, property owner information, and any data derived from or related to the Torrens title register.
Personal Information (PII)	Customer or employee personally identifiable information including names, addresses, identification numbers, contact details, financial status.
Commercially Sensitive Information	Internal business strategies, pricing, tender documents, contracts, partner agreements, or any information that could disadvantage LRS if disclosed.
Financial Data	Budget figures, revenue data, financial forecasts, account details, transaction records, passwords, or credentials.
Regulated Data	Data subject to specific regulatory obligations including court orders, government data sharing agreements, or law enforcement requests.
Proprietary Technical Assets	Source code, system architecture, internal algorithms, database schemas, configuration files, or security control details.

Internal Communications	Internal emails, meeting notes, board papers, executive communications, or any internal discussion not intended for public release.
P&C and Employee Data	Performance records, disciplinary records, remuneration details, medical information, or any data covered by employment agreements.

8 GUARDRAILS FOR ALL AI USE

The following guardrails apply to all use of AI tools by LRS staff, regardless of tool category:

De-Identification Before Input - Always remove or anonymise any personal, internal, or sensitive identifiers (names, IDs, financial figures, case references, addresses, property details) before inputting data into any Category B AI tool. Be aware that removing a name but retaining a specific address, title reference, or dealing number may still be identifying. When in doubt, do not proceed - consult the IT Service Desk.

Verification Before Use - Always verify AI-generated content for accuracy and bias before using it in any official capacity. The level of verification required is proportionate to risk (see Rule 2 in Section 4). The fact that AI produced a plausible-sounding output is not evidence that it is accurate.

No Verbatim Copying Without Attribution - Do not copy AI-generated content verbatim into official LRS documents, correspondence, or records without clearly identifying it as AI-generated and subject to human review. AI-generated content that has been reviewed and adopted by a human should be attributed appropriately.

No Sole Reliance on AI for Recommendations - AI output must never be the sole basis for a formal recommendation, decision, or official record. It may inform and assist human judgement but not replace it.

Record Retention of AI Interactions - For any AI interaction that contributes to an official output (a report, recommendation, formal correspondence, or decision), retain a record of the prompt(s) and key outputs for audit purposes. Retention requirements align with the relevant record type under LRS's records management policy.

No Uploading LRS Documents or Screenshots to Unapproved Tools - Do not upload LRS documents, screenshots of LRS systems, or images containing LRS information to any Category B or Category C AI tool, even for summarisation or analysis purposes.

Intellectual Property Awareness - Be aware that AI-generated content may reproduce content protected by third-party copyright. Do not use AI-generated text, images, or code in LRS outputs without considering whether copyright obligations apply. Seek guidance from Legal if uncertain.

9 AI FEATURES IN ENTERPRISE SOFTWARE

Many enterprise software platforms used by LRS include AI features that may be activated by default, through configuration, or through user action. These features are not automatically approved for use simply because the underlying platform is approved.

The AISC maintains a register of enterprise AI features and their approval status. Before using any AI feature in an enterprise platform, check the AISC register or contact the IT Service Desk to confirm whether the feature is:

- Approved: Can be used within specified parameters
- Pending Review: Do not use until AISC decision.
- Restricted: May only be used under specified conditions
- Prohibited: Must not be used; contact IT if the feature is active

Where an enterprise AI feature processes or has access to LRS data (including email, documents, calendar items, or system records), it must be treated as a Category A system and is subject to all associated requirements including data classification, human oversight, and audit logging.

10 COMPLIANCE, REPORTING AND MONITORING

10.1 MANDATORY TRAINING AND ATTESTATION

All staff must complete mandatory AI Literacy training and provide a written attestation that they have read and understood this policy before using any AI tool for LRS work purposes. Training must be repeated annually.

LRS has an existing training provider with an AI literacy module that can be used for this requirement. Upon completion of the training module, staff must attest in writing (via the training platform) that they have:

Completed the AI literacy training module.

Read and understood the LRS Generative AI End User Policy and associated elements of the AI Governance Framework.

Agreed to comply with its requirements.

Where the training module is not yet accessible to all staff, the following interim measures apply:

All staff must read and sign this policy before using any AI tool.
Managers must provide team briefings on AI policy obligations.
High-risk AI use (Tier 3 or Tier 4 risk) requires supervisor sign-off until formal training is complete.

10.2 INCIDENT REPORTING

Staff who become aware of a potential data breach, policy violation, or AI-related incident must report it immediately to their manager and the IT Security Team, in accordance with the LRS Data Breach Response Plan. Do not delay reporting to investigate the matter yourself.

AI incidents are classified into the following severity levels:

Severity	Examples	Response
Low	AI generates a biased or inaccurate output that is caught during human review and not acted upon; staff member uses a prohibited tool for a low-risk task.	Log the incident. Manager notified. Refresher guidance provided.
Moderate	Staff member pastes internal (non-sensitive) content into a Category B tool; AI output used without human verification in a non-critical context.	IT Security Team notified within 24 hours. Incident recorded. Manager and CIO informed.
High	Sensitive or Internal data submitted to an unapproved AI tool; AI output used without verification in an official record	IT Security notified immediately. AISC notified within 4 hours. Data Breach Response Plan activated if applicable.

	or customer communication; suspected prompt injection attack.	
Critical	Data breach exposing PII or land titles data through AI tool; AI output used to make an irreversible decision affecting individual rights; suspected malicious use or AI-assisted fraud.	CEO and AISC notified immediately. Data Breach Response Plan activated. External reporting obligations assessed. Incident recorded and preserved for investigation.

10.3 MONITORING

GenAI usage on LRS corporate devices and platforms is actively monitored for traceability and policy compliance. LRS uses Microsoft Purview as the primary platform for monitoring AI tool usage across corporate devices and M365. Monitoring includes:

- Logging of AI tool access on corporate devices (tool name, timestamp, user ID) - via Microsoft Purview or other logging and monitoring capabilities
- Detection and flagging of policy violations (e.g. sensitive data transmitted to blocked AI services)
- Where technically feasible: logging of prompt categories and output types for approved enterprise AI tools
- Log retention period: 6 months.
- Review process: IT Security Team reviews anomalies flagged by Purview; AISC reviews summary logs quarterly.

Staff are notified that AI usage is monitored. Monitoring is proportionate and conducted in compliance with applicable NSW workplace surveillance legislation and LRS employment agreements. The monitoring programme has been reviewed by Legal and P&C.

10.4 COMPLIANCE WITH EXISTING POLICIES

All staff must comply with existing LRS policies when using GenAI, including (but not limited to): the Acceptable Use of IT Policy; employee agreements and codes of conduct; the Privacy Management Plan; and the Records Management Policy. This policy supplements those policies and does not replace them.

11 SANCTIONS AND ENFORCEMENT

Violations of this policy are governed by LRS's existing employment contracts and conduct frameworks.

The AISC may also impose technical access restrictions as a risk management measure, separate from any disciplinary process.

12 THIRD-PARTY SUPPLIERS AND CONTRACTORS

Third-party suppliers and contractors who perform work on behalf of LRS are expected to comply with this policy when performing that work. LRS's approach to third-party AI governance is as follows:

Written Approval Required

Any third-party supplier intending to use AI tools in the delivery of services to LRS must seek written approval from LRS prior to doing so. This ensures LRS can assess whether the proposed use is consistent with this policy, the AI Governance Framework, and applicable concession obligations.

MSA and Contract Requirements

Third party AI use will be covered under Master Service Agreements (MSAs) and Statements of Work (SOWs).

NSW Data Sovereignty - Systems and People

All data sovereignty and NSW-hosted compute requirements apply equally to third-party suppliers. Suppliers cannot route LRS data through AI systems hosted or processed outside NSW. This obligation extends to people: staff and contractors engaged by a supplier to perform AI-related work on LRS data must be physically located within NSW. Offshore or interstate resourcing of AI-related work requires written approval from the LRS CEO.

13 DOCUMENT REVIEW

This policy will be reviewed at least every six months to ensure it remains appropriate and fit for purpose, or sooner in response to any of the following trigger events:

- Significant changes to the AI technology landscape or the LRS AI strategy
- New or amended legislation or regulation affecting AI use.
- Material changes to the organisation's risk profile
- Lessons learned from AI incidents, audits, or external assessments.
- Changes to the list of approved or permitted AI tools.
- Direction from the AISC or the Board

The AI Steering Committee is responsible for managing the review process and obtaining approval for any updates. All changes will be communicated to staff, and training materials will be updated accordingly.

This document is recorded in the ICT Document Library. The responsible division is ICT; the contact is the Chief Information Officer.

14 DOCUMENT VERSION CONTROL

Version	Date	Author	Summary of changes
0.1	11/03/2026	Craig Tidmarsh	Initial draft.
0.2	07/04/2026	Propel Ventures	Second draft incorporating Propel Ventures independent review and consultant-added best-practice improvements. See change log below.
0.3	09/04/2026	Propel Ventures	Third draft incorporating decisions from CEO/Head of Technology interview (9 April 2026). Key changes: Human-in-the-loop and No Automated Decision-Making prominently highlighted; Skunkworks AI call-out added; current AI tool status (Copilot DCR pending, Purview monitoring) documented; training attestation mechanism specified; third-party supplier section added; data governance gap note added; risk tier 'Extreme' renamed 'Critical'; sanctions linked to existing employment contracts and Acceptable Use of IT Systems policy.
0.4	13/04/2026	Propel Ventures	Formatting and terminology updates following internal review. Table of Contents added. Core Principles, Prohibited Activities, and Guardrails sections all reformatted from heading/paragraph style to inline format for readability. Enterprise software examples updated — ServiceNow Now Assist and Salesforce Einstein removed; Adobe Acrobat and Google Chrome AI mode added as more relevant examples. 'HR and Employee Data'

			category renamed to 'P&C and Employee Data' in the data classification table. 'CIO's office' simplified to 'CIO' throughout. Training attestation expanded to include confirmation that staff have read associated elements of the AI Governance Framework, not just the policy. Monitoring logging note broadened from Microsoft Purview specifically to 'Microsoft Purview or other logging and monitoring capabilities'. Monitoring programme sign-off updated from 'Legal and HR' to 'Legal and P&C'. Third-party written approval route updated from 'CIO's office or Procurement team' to 'CIO or Legal team'.
0.5	15/04/2026	Propel Ventures	Version and date updated. Change log heading updated. Minor page number adjustments in Table of Contents. No substantive content changes
0.6	20/04/2026	Craig Tidmarsh	Incorporating feedback from CEO and CIO
1.0	21/04/2026	Mark Hansen	Approve Final revision
1.1	14/05/2026	Mark Hansen Ellen Locke	Incorporation of Board feedback