

AI USE-CASE ASSESSMENT

Community library catalogue assistant

This report combines 4 core assessment results and 1 separately saved policy or guidance review.

Preliminary assessment based on the confirmed use case and customer-provided information. No supporting evidence was supplied or independently verified.

STATUS	REPORT BASIS	SUPPORTING FILES	REPORT VERSION	COUNTRY	PREPARED
Complete	Preliminary		2	AU	21 September 2026

SUGGESTED ROUTE Suitable to continue design with the stated conditions.

This report supports an accountable human decision. It is not legal advice, certification, approval or proof of compliance.

01 / CONFIRMED INPUT

Confirmed input

Working title	Community library catalogue assistant
Business context	Collection Services at a fictional Australian metropolitan public library. Six librarians catalogue about 900 new or donated items each month for a catalogue used by library members and staff.
Current stage	Design
Problem or opportunity	Librarians manually draft plain-language summaries, choose subject headings and check catalogue metadata. Variable wording and backlogs can delay when new items become discoverable.
What happens today	A librarian checks the supplier record and the physical item, drafts a summary, selects controlled subject headings, enters metadata and asks a senior cataloguer to review unusual records. The median backlog is described as 12 working days, but no supporting record has been supplied.
Evidence available	The use-case owner reports a 12-working-day median backlog and recurring rework for inconsistent summaries. These are stakeholder claims in this workbook; no supporting evidence is uploaded for this baseline run.
Outcome sought	Make new items discoverable sooner while preserving accurate, inclusive and consistent catalogue descriptions.
How success will be measured	During a supervised pilot, compare handling time, correction rate, prohibited or outdated subject headings, accessibility of summaries and the number of records rejected by librarians. No production target has an approval claim that requires verification.
What AI will do	Suggest a short catalogue summary, candidate controlled subject headings and possible missing metadata fields for a librarian to review.
Users	Cataloguing librarians and senior cataloguers would use the assistant. Library members, authors and communities described by catalogue terms could be affected by inaccurate, exclusionary or misleading metadata.
Affected people	Cataloguing librarians and senior cataloguers would use the assistant. Library members, authors and communities described by catalogue terms could be affected by inaccurate, exclusionary or misleading metadata.
Owner and next decision-maker	The Manager, Collection Services owns the outcome and may approve a sandbox evaluation. The Library Governance Committee decides whether any supervised pilot can begin.

Decision boundary	When uncertain, the assistant must flag the field or suggestion for human review rather than inventing or publishing an answer. A cataloguing librarian must check the source record, wording, subject headings and missing fields and may edit or reject every suggestion. Contested or sensitive descriptions must be escalated to a senior cataloguer. Nothing may be written to the live catalogue during the sandbox.
Information used	Bibliographic supplier records, publisher descriptions, the physical-item record, the library's approved subject-heading vocabulary and local cataloguing guidance. The sandbox must use fictional or public-domain examples and must not include member borrowing records.
Output and resulting action	The assistant produces a draft summary, suggested headings and missing-field prompts. A librarian may edit or reject every suggestion. Nothing is written to the live catalogue during sandbox testing.
Human responsibility	A cataloguing librarian checks the source record, wording, subject headings and missing fields and may reject the draft. A senior cataloguer reviews contested or sensitive descriptions. The Manager may pause testing, and overrides and reasons must be recorded.
Initial scope and exclusions	A six-week sandbox using 150 fictional or public-domain item records. Exclude live catalogue writes, member data, automated publishing, acquisitions decisions, censorship decisions and changes to the controlled vocabulary.
Other ways to address the problem	The library could improve templates, training, queue allocation and deterministic vocabulary lookup without generative AI. AI may help draft summaries, but the sandbox should compare it with those simpler options.
If it goes wrong	A misleading summary or subject heading could make an item hard to find, misrepresent a person or community, or introduce inappropriate language. If a librarian misses the error, it could enter a pilot dataset and be repeated. Sandbox isolation and mandatory review should interrupt the chain before publication.
Dependencies and constraints	The approved vocabulary, current cataloguing guidance, a versioned test set, trained reviewers, an override log and a named escalation owner must be available. Testing must record the model version and prevent live catalogue integration.
Proposed next step	The Manager, Collection Services should lead the sandbox design with cataloguers, accessibility advisers, community engagement, privacy and technology staff. The Library Governance Committee should review the test plan before testing begins.

02 / CORE ASSESSMENT

Core assessment

01 Business assessment

Information strength: Limited

Report basis: Preliminary - customer-provided information; no supporting evidence.

Sandbox design is clear, but baseline evidence is thin

The proposed sandbox is well bounded and has a clear human review step. The main gap is that the backlog and rework baseline are still customer claims, so the business case is only partly evidenced.

Findings

- The use case has a clear design-stage problem, outcome and workflow for a catalogue assistant.
- The reported 12-working-day backlog and recurring rework are still customer claims and have no supporting baseline record.
- The sandbox is bounded to fictional or public-domain records and excludes live catalogue writes and member data.
- Funding, staffing effort and reviewer capacity for the sandbox are not set out.
- The use case names sensible alternatives such as templates, training, queue allocation and deterministic vocabulary lookup, but it does not compare them in a documented way yet.

Information used

- Confirmed use-case version 1

Limitations

- No supporting evidence record was supplied for the reported backlog baseline.
- This assessment is limited to the business case and does not judge technical, legal, privacy or policy matters.
- The pilot measures are future-state measures and should not be treated as current results.

Review reference: REV-35B01A9E

02 People assessment

Information strength: Limited

Report basis: Preliminary - customer-provided information; no supporting evidence.

Sandbox design is workable, with a few people gaps

The use case is ready for sandbox design in people terms, with clear users, human review and a narrow test boundary. A few consultation, accessibility and training details still need to be confirmed before testing starts.

Findings

- The sandbox has a clear human review chain, with librarian review, senior cataloguer escalation and manager pause rights.
- The review rules do not yet define the unacceptable language and accuracy thresholds that would trigger rejection or escalation.
- The use case says accessibility matters, but it does not yet set out specific accessibility needs, checks or success criteria for the sandbox.
- The plan names the right internal roles for the next step, but there is no record of consultation with cataloguers, accessibility advisers or community engagement staff.
- The proposed workflow relies on trained reviewers and an override log, but the training scope and workload plan are still open.
- The sandbox boundary is narrow enough for design-stage testing because it excludes live writes, member data and automated publishing.

Information used

- Confirmed use-case version 1

Limitations

- No supporting evidence record was supplied for the alleged backlog or rework baseline, so that claim is treated as customer-provided information only.
- The assessment is limited to people and work impacts at design stage and does not assess business value, technical performance, privacy, legal compliance or policy compliance.
- Consultation, accessibility checks and reviewer readiness are proposed next-step items, not already confirmed actions.

Review reference: REV-940DBD9C

03 Technology assessment

Information strength: Limited

Report basis: Preliminary - customer-provided information; no supporting evidence.

Sandbox design is ready, with a few missing test details

The proposed library catalogue assistant is technically workable for a controlled sandbox at the current design stage. The main gap is that the test set, thresholds and stop rules still need to be set before testing begins.

Findings

- The proposed task is bounded enough to design a controlled sandbox for catalogue drafting support.
- The sandbox still needs agreed thresholds, a representative test set and clear stop rules.
- The current evidence does not show how retention, supplier terms and test-set provenance will be handled.
- Reviewer workload and override logging are named, but they are not yet tested in practice.
- The use case defines a clear human decision boundary and forbids live catalogue writes during sandbox testing.

Information used

- Confirmed use-case version 1

Limitations

- No supporting evidence records were supplied beyond the customer-provided use case, so this is a preliminary design-stage assessment.
- Operational controls, thresholds and data-handling details remain unverified because they are listed as unknowns in the use case.

Review reference: REV-C050AE07

04 Risks and oversight

Information strength: Limited

Report basis: Preliminary - customer-provided information; no supporting evidence.

Sandbox can be designed, but key governance gaps remain

The use case is bounded enough to keep working on a sandbox design. But the review boundary, stop rules and escalation chain still need to be made explicit before testing starts.

Findings

- The proposed task is narrow enough for design-stage review because it only drafts summaries, headings and missing-field prompts for librarian review.
- A misleading or exclusionary heading could harm discoverability or misrepresent people and communities if review misses the error.
- The sandbox stop criteria are listed as unknown and are not yet set out in a way that can be checked at the gate.
- The accountable escalation owner for testing is mentioned as a dependency, but the operational chain for pause, review and override is not fully pinned down.
- Baseline backlog, rework rates, thresholds for unacceptable language and reviewer workload are still unknown.

Information used

- Confirmed use-case version 1
- Customer starting input — Community-Library-Catalogue-Starter-Pack-v0.1.xlsx

Limitations

- This is a design-stage assessment, so it does not test whether controls work in practice.
- The available evidence is limited to the confirmed use case and the supplied starter input; no independent operational records were provided.
- The recorded backlog and rework figures are customer claims, not independently verified evidence.

Review reference: REV-1F20AE69

03 / POLICY AND GUIDANCE

Policy and guidance reviews

P1

Guidance for AI adoption: foundations

Voluntary public guidance

National AI Centre - Published 5 May 2026 - Australia

Scope: Organisations starting to use AI, using it in lower-risk ways, or establishing basic responsible-AI governance in a business context.

Report basis: Preliminary - customer-provided information; no supporting evidence.

Information strength: Limited

Sandbox is workable, but key test details are still missing

The use case is well bounded for a design-stage sandbox and has a clear human review chain. The main gap is that stop rules, thresholds, test-set provenance and workload details are still unresolved.

Review status: No stop

Findings

- The guidance is applicable to this use case because it is a design-stage, lower-risk sandbox for an AI drafting assistant in a public library setting.
- The sandbox stop rules, thresholds for unacceptable language and accuracy, and test-set provenance are still unresolved.
- The use case describes a human review chain with librarian review, senior cataloguer escalation and manager pause rights.
- A misleading or exclusionary heading could harm discoverability or misrepresent people and communities if review misses the error.
- The evidence does not yet show how reviewer workload, training scope, retention and supplier terms will be managed.

Overlaps

- The business, people, technology and risk assessments all agree the sandbox is bounded and suitable for design-stage work.
- All four core assessments identify missing thresholds, stop rules or governance details as the main gap.
- The confirmed use case and starter input repeat the same accountable owner, review chain and sandbox exclusions.

Conflicts for human review

- The core risk assessment notes the baseline backlog and rework figures are customer claims, which means they should not be treated as independent evidence.
- The policy review treats the source package as preliminary because the evidence base is limited to customer-provided material and a single official guidance source.

Limitations

- This is a preliminary evidence review, not legal advice or certification.
- The backlog and rework baseline remain customer claims with no supporting baseline record.
- No independent operational records were supplied for reviewer workload, test thresholds or retention handling.

Authoritative source practices

- Decide who is accountable: <https://www.ai.gov.au/staying-safe-and-responsible/essential-ai-practices/guidance-ai-adoption-foundations#1-decide-who-is-accountable>
- Maintain human control: <https://www.ai.gov.au/staying-safe-and-responsible/essential-ai-practices/guidance-ai-adoption-foundations#6-maintain-human-control>
- Understand impacts and plan accordingly: <https://www.ai.gov.au/staying-safe-and-responsible/essential-ai-practices/guidance-ai-adoption-foundations#2-understand-impacts-and-plan-accordingly>
- Measure and manage risks: <https://www.ai.gov.au/staying-safe-and-responsible/essential-ai-practices/guidance-ai-adoption-foundations#3-measure-and-manage-risks>
- Test and monitor: <https://www.ai.gov.au/staying-safe-and-responsible/essential-ai-practices/guidance-ai-adoption-foundations#5-test-and-monitor>
- Share essential information: <https://www.ai.gov.au/staying-safe-and-responsible/essential-ai-practices/guidance-ai-adoption-foundations#4-share-essential-information>

Source captured 13 September 2026. Review reference: REV-A77391B4.

04 / DECISION RECORD

Conditions, gaps and limitations

Conditions

- The use case may affect communities through wording and subject headings, so human review and escalation are important.
- The reported backlog and rework are stakeholder claims and still need supporting evidence before any performance claim is relied on.
- This is a sandbox design, not approval to deploy or publish.
- Confirm reviewer training scope, workload assumptions, retention handling and supplier terms for the sandbox. Owner: Manager, Collection Services. Due: Before testing begins. Closure evidence: Signed-off operating plan covering training, workload, retention and supplier terms.
- Document the sandbox stop rules, the rejection and escalation thresholds, and the provenance and selection method for the test set. Owner: Manager, Collection Services. Due: Before testing begins. Closure evidence: Approved sandbox test plan with named thresholds, stop rules and test-set record.
- Set explicit rules for unacceptable language, accuracy checks and escalation of contested or sensitive descriptions. Owner: Senior cataloguer. Due: Before testing begins. Closure evidence: Written review rubric and escalation guide for cataloguers.
- Write the human review and stop rules for sandbox testing, including when the librarian must escalate contested or sensitive descriptions and when testing must pause. Owner: Manager, Collection Services. Due: Before sandbox testing begins. Closure evidence: Approved sandbox governance pack with the review boundary, escalation triggers and stop criteria recorded..
- Name one accountable escalation owner for test incidents and define the pause, review and override chain. Owner: Manager, Collection Services. Due: Before sandbox testing begins. Closure evidence: Escalation register or RACI that names the owner and the pause path..
- Obtain, review and approve the information needed to resolve: A documented review rubric for unacceptable language, accuracy and contested or sensitive descriptions. Owner: Owner to be assigned. Due: Before the next governance decision. Closure evidence: An approved record demonstrating that a documented review rubric for unacceptable language, accuracy and contested or sensitive descriptions is resolved..

Related evidence gaps

Testing, monitoring and evidence

- The reported 12-working-day backlog and recurring rework are still customer claims and have no supporting baseline record.
- The review rules do not yet define the unacceptable language and accuracy thresholds that would trigger rejection or escalation. The sandbox stop criteria are listed as unknown and are not yet set out in a way that can be checked at the gate. The sandbox stop rules, thresholds for unacceptable language and accuracy, and test-set provenance are still unresolved.
- Defined accessibility checks for draft summaries and subject headings, including the people who will run them and the records they will keep. A record of consultation with cataloguers, accessibility advisers and community engagement staff on the sandbox design. A documented review rubric for unacceptable language, accuracy and contested or sensitive descriptions.
- The sandbox still needs agreed thresholds, a representative test set and clear stop rules.

- Representative test-set definition and provenance for the 150 fictional or public-domain item records.; Retention period, supplier term and local handling rules for sandbox records.

Contributing saved outputs: Business assessment; People assessment; Technology assessment; Risks and oversight; Guidance for AI adoption: foundations

Data and information

- A dated baseline record for backlog and rework, with the source extract used to calculate it.

Contributing saved outputs: Business assessment

Accountability and human decisions

- A simple resourcing note for reviewer time, sandbox support and any funding needed.
- A short comparison method for the AI draft approach and at least one non-AI alternative. Baseline measures for backlog, rework, reviewer workload and the threshold for unacceptable language. A written plan for reviewer training, workload allocation, retention handling and supplier terms.
- Written reviewer thresholds for unacceptable language, inaccurate summaries and subject-heading errors, including reject, edit and escalation triggers. A written sandbox stop rule with concrete stop triggers, pause authority and review steps. A signed sandbox test plan with stop rules, rejection thresholds, escalation rules and test-set provenance.
- Reviewer training scope, workload allocation and the override log procedure for the sandbox. Approved sandbox acceptance thresholds, including acceptable error rates, correction rates and stop conditions.; Representative test-set definition and provenance for the 150 fictional or public-domain item records.; Documented stop rules and escalation path for unsafe or poor-quality outputs.
- The accountable escalation owner for testing is mentioned as a dependency, but the operational chain for pause, review and override is not fully pinned down.
- A named escalation owner and a simple incident path for pause, review and override decisions.

Contributing saved outputs: Business assessment; People assessment; Technology assessment; Risks and oversight; Guidance for AI adoption: foundations

People and use

- The plan names the right internal roles for the next step, but there is no record of consultation with cataloguers, accessibility advisers or community engagement staff.

Contributing saved outputs: People assessment; Guidance for AI adoption: foundations

Limitations

- No supporting evidence record was supplied for the alleged backlog or rework baseline, so that claim is treated as customer-provided information only.
- The pilot measures are future-state measures and should not be treated as current results.
- Consultation, accessibility checks and reviewer readiness are proposed next-step items, not already confirmed actions.
- No independent operational records were supplied for reviewer workload, test thresholds or retention handling.
- This report supports human review. It is not legal advice, certification, approval or proof of compliance.

Saved output index

Output	Version or result hash	Status
Confirmed AI use case	Version 1	confirmed
Business assessment	35b01a9eb2e213d97a942395a6936b780b851992ddb58405946851194f32c33e	viewed
People assessment	940dbd9c76eb59743af7cfb0acb8558e20711c6805dc56f768f3e1a5e883cba4	viewed
Technology assessment	c050ae0716c3b1bae077458e3f8692323fb25c06c5f3ef291520f2d600e11571	viewed
Risks and oversight	1f20ae69e46005f22597f7dd36efaa9df00efb32832cc56a7fb8aa20a615ee94	viewed
Guidance for AI adoption: foundations	Published 5 May 2026 · a77391b4969c8df6992b4e24d169a1afeb55646f1561245300ae8cf2697be858	viewed

Report identity: 831f531b3f4dacebf70ea60687d64c1157866943b0c218228f24caf9593edb91

Confirmed use-case version: 1. Source assessment revision: 22.

Always consult the linked authoritative policy sources. This report is designed for reading, searching and copying.