

AI USE-CASE ASSESSMENT

Facilities maintenance request triage assistant

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

Evidence-backed assessment using 1 supporting file. Findings can be traced to those files, but this report does not independently verify their accuracy or prove readiness.

STATUS Complete	REPORT BASIS Evidence-backed	SUPPORTING FILES 1	REPORT VERSION 2	COUNTRY AU	PREPARED 21 September 2026
--------------------	---------------------------------	-----------------------	---------------------	---------------	-------------------------------

SUGGESTED ROUTE Suitable to enter controlled sandbox testing 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	Facilities maintenance request triage assistant
Business context	A fictional Australian property-services team is designing an AI assistant for building-maintenance requests.
Current stage	Design stage, with a controlled pilot proposed.
Problem or opportunity	The team needs a way to classify maintenance requests, suggest urgency and routing, and draft a work order so coordinators can handle requests faster and more consistently.
What happens today	Emergency reports bypass the assistant and go straight to the duty coordinator. For other requests, a facilities coordinator reviews the original request and building details, then may change, reject or escalate every suggestion before scheduling.
Evidence available	An approved synthetic baseline and proposed pilot controls are available in the supporting evidence. They cover the current process baseline, operating boundary, human oversight, pilot thresholds, stop rules and evidence still required after the pilot.
Outcome sought	A controlled pilot-ready design for a maintenance request triage assistant.
How success will be measured	The pilot thresholds in the evidence are 100% recall for safety-critical examples with zero critical false negatives, at least 95% correct routing for routine requests, recorded coordinator correction and override rates by category and severity, logged escalation and bypass events, and no autonomous scheduling or dispatch event.
What AI will do	Classify maintenance requests, suggest urgency and routing, and draft a work order.
Users	Facilities coordinators and the duty coordinator.
Affected people	People who submit building-maintenance requests, building occupants, contractors and the facilities team.
Owner and next decision-maker	Facilities Operations Manager.
Decision boundary	The AI may suggest classification, urgency, routing and a draft work order. A human must inspect the original request and building details, then may change, reject or escalate every suggestion before scheduling. The AI must not schedule work, dispatch a contractor, close a request or make a safety decision.
Information used	Original request details, building details, affected asset, proposed category, urgency and reason.

Output and resulting action	A classified request, suggested urgency, suggested routing and a draft work order for human review.
Human responsibility	A facilities coordinator must inspect the request and building details, review the original request, building, affected asset, proposed category, urgency and reason, and decide whether to change, reject, escalate or accept the suggestions before scheduling. Overrides and reasons must be recorded.
Initial scope and exclusions	In scope: non-emergency maintenance requests. Exclusions: emergency reports, scheduling work, contractor dispatch, closing requests and safety decisions.
Other ways to address the problem	Emergency reports go directly to the duty coordinator.
If it goes wrong	If the assistant gives a poor suggestion, a human review step still sits before scheduling, but delays or misrouting could affect maintenance handling. The pilot must stop immediately after any missed emergency, any unauthorised scheduling or dispatch, or two material misroutes within a rolling five-day period.
Dependencies and constraints	Controlled pilot only; emergency reports bypass the assistant; the Facilities Operations Manager approves any pilot; the duty coordinator may pause processing immediately; a coordinator may review no more than 40 AI-triaged items in one shift during the pilot.
Proposed next step	Prepare for a controlled pilot.

02 / CORE ASSESSMENT

Core assessment

01 Business assessment

Evidence strength: Limited

Report basis: Evidence-backed - 1 cited supporting file; accuracy is not independently verified.

Pilot-ready design is defined, but not yet proven

The design and controlled pilot boundary are well specified for a non-emergency maintenance triage use case. The main gap is that the pilot controls and thresholds are only proposed, so the next step is to get formal approval to run the supervised pilot.

Findings

- The use case has a clear business problem and a defined pilot outcome: triage non-emergency maintenance requests more quickly and consistently.
- The business has set measurable pilot thresholds, including safety-critical recall, routine routing accuracy, override tracking and stop rules.
- The operating model keeps a human in charge before scheduling, with the coordinator able to change, reject or escalate every suggestion.
- No funding, staffing or time estimate is provided for the pilot.
- The business case does not compare the pilot against other non-AI process options in a structured way.
- The pilot controls are described as approved synthetic baseline and proposed controls, but this assessment cannot independently verify that they have been signed off as operationally ready.

Evidence used

- Confirmed use-case version 1
- Facilities-Maintenance-Triage-Evidence.md

Limitations

- The assessment is limited to the supplied confirmed use case and one supporting evidence record.
- The evidence is synthetic and fictional, so this assessment judges business readiness for the stated stage rather than real-world performance.
- No independent verification of approval, implementation or pilot operation was available in the supplied material.

Review reference: REV-826493CB

02 People assessment

Evidence strength: Limited

Report basis: Evidence-backed - 1 cited supporting file; accuracy is not independently verified.

Pilot-ready design has a clear human review step

The proposed design keeps a person in charge of triage decisions and sets a defined pilot boundary. The main gap is missing consultation and workforce-readiness detail for the people who will use and handle the process.

Findings

- The proposed workflow keeps the coordinator as the decision-maker for non-emergency requests.
- There is no recorded consultation or training plan for the staff who will use the assistant.
- Accessibility needs are not described beyond the information that must be shown to the coordinator.
- The pilot will record overrides, corrections and escalation events, but no direct experience measure is set for ease of use or burden.
- The people who submit requests and those who handle them may face delays or confusion if the triage messages are unclear.

Evidence used

- Confirmed use-case version 1
- Facilities-Maintenance-Triage-Evidence.md

Limitations

- This review uses only the confirmed use case and supplied supporting evidence.
- The supporting evidence is synthetic and does not show live use or consultation outcomes.
- No prior assessment stage was used as a source.

Review reference: REV-CD8716F9

03 Technology assessment

Evidence strength: Limited

Report basis: Evidence-backed - 1 cited supporting file; accuracy is not independently verified.

Pilot-ready design is broadly defined, with evidence gaps for testing

The proposed pilot has a clear boundary, human review step and stop rules. The main gap is that the actual pilot results do not yet exist, so production readiness cannot be shown.

Findings

- The proposed pilot has a clear decision boundary and human review step before scheduling.
- Emergency reports bypass the assistant and go straight to the duty coordinator.
- Pilot effectiveness has not yet been demonstrated with live or labelled pilot results.
- The evidence still needed after the pilot includes monitoring, incident and remediation records.
- Current access and retention controls are not yet described in enough operational detail.

Evidence used

- Confirmed use-case version 1
- Facilities-Maintenance-Triage-Evidence.md

Limitations

- The supporting evidence is synthetic and describes proposed controls, not proven performance.
- No live pilot results were supplied, so production readiness cannot be assessed.
- This assessment is limited to technical and operational feasibility.

Review reference: REV-53D771F5

04 Risks and oversight

Evidence strength: Limited

Report basis: Evidence-backed - 1 cited supporting file; accuracy is not independently verified.

Pilot design is bounded, with one key oversight gap

The proposed controlled pilot has a clear boundary, human review and stop rules. The main gap is that escalation and challenge routes exist in the process, but the record does not yet show a separate way for affected people to contest outcomes.

Findings

- The pilot has a clear non-emergency boundary and keeps safety decisions, scheduling and dispatch out of the AI's hands.
- The process still leaves room for delay or misrouting if the AI suggestion is wrong, even though a human review step remains.
- The record does not describe a separate contest path for affected requesters beyond the coordinator's internal review.
- Accountable authority is named, and stop and pause powers are assigned to specific roles.
- The evidence shows the proposed controls on paper, but it does not yet show that they work in live use.

Evidence used

- Confirmed use-case version 1
- Facilities-Maintenance-Triage-Evidence.md

Limitations

- The assessment relies on a fictional synthetic baseline and proposed pilot controls, not live operational evidence.
- The sources do not provide a separate requester contest route, so that issue is an evidence gap rather than a proven failure.
- This review only covers risks, oversight and control design at the stated design-stage pilot step.

Review reference: REV-E164AA50

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: Evidence-backed - 1 cited supporting file; accuracy is not independently verified.

Evidence strength: Limited

Pilot design is clear, but evidence gaps remain

The proposed guidance fit is generally good for a controlled pilot of non-emergency maintenance triage. The main gaps are that the controls are still proposed, and live proof of testing, monitoring and remedy is not yet available.

Review status: No stop

Findings

- The use case sits within the guidance scope for a controlled pilot in a business setting.
- The process keeps a human in charge before scheduling, and emergency reports bypass the assistant.
- The pilot controls are described as proposed and synthetic, but the supplied material does not show live implementation or proof of effectiveness.
- The material does not show a separate contest path for affected people beyond internal coordinator review.
- Accountability is named for the business owner, the duty coordinator and the coordinator review step.

Overlaps

- All four core assessments agree that the pilot boundary is clear and that human review stays in place before scheduling.
- All four core assessments agree that the main gap is missing proof of live or complete pilot evidence.

Conflicts for human review

- The risks review highlights a missing separate contest path for affected people, while the people review focuses more on consultation and workforce-readiness gaps.

Limitations

- The review is limited to the supplied confirmed use case, one supporting evidence record and four core assessment results.
- The evidence is synthetic and fictional, so the review can only judge the stated pilot design, not live performance.
- No independent proof of implementation, approval or operating results was supplied.

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-CF413ABE.

04 / DECISION RECORD

Conditions, gaps and limitations

Conditions

- This is a higher-impact operational workflow because it affects maintenance routing and urgency.
- The pilot controls rely on the supplied evidence and have not been independently verified here.
- Approve the controlled pilot plan, including the operating boundary, stop rules, measures and evidence pack. Owner: Facilities Operations Manager. Due: Pilot launch gate. Closure evidence: Signed pilot approval or equivalent formal authorisation record..
- Record a short consultation and training plan for facilities coordinators and the duty coordinator, including who is consulted, what they are told and how they will be trained before pilot start. Owner: Facilities Operations Manager. Due: Before controlled pilot start. Closure evidence: Approved pilot brief or training plan showing consultation topics, attendees and training material..
- Run the controlled pilot with the labelled test set and record recall, routing accuracy, override rates, escalation and bypass events, and any stop-rule trigger. Owner: Facilities Operations Manager. Due: Before any production decision. Closure evidence: Signed pilot test report with the labelled result set, confusion matrix, override log and stop-rule record.
- Run the supervised pilot and collect the missing evidence on labelled test results, monitoring, incidents, overrides and remediation records. Owner: Facilities Operations Manager. Due: Before any broader rollout decision. Closure evidence: Signed pilot report with test results, monitoring log, incident log, override log and remediation record.

Related evidence gaps

People and use

- No funding, staffing or time estimate is provided for the pilot.
- Pilot budget, staffing load and coordinator time allocation.

Contributing saved outputs: Business assessment

Other evidence gaps

- The business case does not compare the pilot against other non-AI process options in a structured way.
- A short comparison of the main non-AI alternative for routine requests and why the AI pilot is preferred.
- Pilot effectiveness has not yet been demonstrated with live or labelled pilot results.
- The pilot controls are described as proposed and synthetic, but the supplied material does not show live implementation or proof of effectiveness.

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

Data and information

- There is no recorded consultation or training plan for the staff who will use the assistant. Accessibility needs are not described beyond the information that must be shown to the coordinator.
- Documented access and retention controls for the pilot data, model and prompt versions.; Monitoring and incident procedure records for pilot operation and post-pilot remediation.

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

Testing, monitoring and evidence

- A brief record of consultation with facilities coordinators and the duty coordinator on the proposed workflow and pilot setup. The planned accessibility supports for the review screen, work order draft and request-handling messages. A documented path for affected people to question or contest outcomes, including response handling.
- The record does not describe a separate contest path for affected requesters beyond the coordinator's internal review.
- A simple measure of coordinator experience, such as clarity, confidence or workload feedback during the pilot. A documented route for requesters or affected people to question, correct or escalate a triage outcome.
- Completed labelled pilot results for the full safety-critical and routine test set, including confusion matrix and recall outcomes. The post-pilot evidence pack named in the supporting evidence, including test results, confusion matrix, workload results and incident procedures.
- A pilot monitoring template that captures correction, override, escalation and bypass trends by category and severity.
- Signed pilot results showing test performance, override rates, incidents and remediation actions.

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

Limitations

- Customer-provided supporting evidence was not independently verified.
- No independent verification of approval, implementation or pilot operation was available in the supplied material.
- No independent proof of implementation, approval or operating results was supplied.
- No prior assessment stage was used as a source.
- 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	826493cb3644de31d21e4e2ab37cea7864068939953bd92dbcf a3084f1717d14	viewed
People assessment	cd8716f96789a9cb14d20f6225fa01c625d58808151c3b682b76 16b736baddcf	viewed
Technology assessment	53d771f5a18bf7d8526c9455345be9d8862da51c714397bfc119 b0a7a43460fb	viewed
Risks and oversight	e164aa50f589b826b1385d7872bcb6c8f7fe74c60130e541c90c 247ec839e27c	viewed
Guidance for AI adoption: foundations	Published 5 May 2026 · cf413abee24484fa86a29e3c8e52604f1 0a89223b4c6a24108a939eea1fb229c	viewed

Report identity: b6e6c1c35c8d375b071e9c1d53a95ebebfa9687b11df7e8a8c5c0499d4d2ea01

Confirmed use-case version: 1. Source assessment revision: 19.

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