

AI USE-CASE ASSESSMENT

Harbour Lane customer-support 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 20 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	Harbour Lane customer-support assistant
Business context	Customer Operations at Harbour Lane Retail, a fictional Australian online retailer. A 12-person support team handles about 4,200 email and web-form enquiries each month.
Current stage	Discovery
Problem or opportunity	Support agents manually classify enquiries, choose a queue, search approved help content and draft replies. Re-routing and repeated searching delay the first response and can make answers inconsistent.
What happens today	An agent opens the message, checks the order system when needed, chooses a category and priority, searches approved help content, drafts a reply and sends it. Team leads handle unusual, sensitive or high-value cases.
Evidence available	Fictional baseline data shows 4,200 enquiries a month, a 19-hour median first response, 23% re-routing, an 11% seven-day reopen rate and 1.6% material policy errors. Mandatory escalations are not measured consistently. No live-customer results are available. E-03 sets draft operating boundaries and oversight controls for the internal evaluation.
Outcome sought	Give customers a faster and more consistent first response without weakening customer care, privacy or human responsibility.
How success will be measured	Reduce median first response from 19 to 12 hours and re-routing from 23% to below 12%. Keep material policy errors below 2%, detect at least 95% of mandatory escalations and do not worsen the seven-day reopen rate.
What AI will do	Classify an enquiry, suggest priority and routing, retrieve approved help content, draft a source-linked reply and flag uncertainty or mandatory escalation for human review.
Users	Support agents and team leads would use the assistant. Customers could be affected, particularly people with limited English, disability, hardship or heightened vulnerability. Support staff would also be affected by changes to their work.
Affected people	Support agents and team leads would use the assistant. Customers could be affected, particularly people with limited English, disability, hardship or heightened vulnerability. Support staff would also be affected by changes to their work.
Owner and next decision-maker	The Head of Customer Operations owns the service and may approve or pause the internal evaluation. A senior AI-governance owner and operational incident owner for any live pilot have not yet been appointed.

Decision boundary	The assistant may classify an enquiry, suggest priority and routing, retrieve approved help content, draft a source-linked reply, and flag uncertainty or mandatory escalation. A support agent must verify the classification, relevant order facts, cited source and draft, and may change or discard it. Team leads handle sensitive or high-risk cases, and authorised people retain all refund, account, fraud and sending decisions.
Information used	The internal evaluation uses synthetic enquiries, separately approved de-identified examples, approved help content and minimum relevant fictional order details. Live credentials, payment data and identity documents are prohibited.
Output and resulting action	The assistant suggests a category, priority, queue, cited source and reply draft. During the internal evaluation, no reply is sent. In any later pilot, a support agent would review, change or discard the draft and decide whether to send it.
Human responsibility	A support agent checks the classification, order facts, cited source and draft, and may change or discard it. Team leads handle hardship, vulnerability, injury, discrimination, threats, fraud, account-security, privacy, legal and uncertain-identity cases. Authorised people make all refund, account and fraud decisions.
Initial scope and exclusions	Run an eight-week internal evaluation using synthetic and approved de-identified cases for delivery, product, return and general order enquiries received by email or web form. Exclude automatic sending, refunds, account changes, fraud, hardship, legal replies, emotion inference and learning from live messages.
Other ways to address the problem	The team could improve queue rules, forms, templates, search and staffing without AI. AI may add value when an agent needs to combine classification, routing, approved source retrieval and a draft, but the evaluation must compare it with these simpler options.
If it goes wrong	A wrong output could delay help, misstate customer rights, expose information or miss a sensitive case. Staff may rely on a weak draft or spend more time correcting it. No output may be sent or used for a high-impact decision without an authorised person.
Dependencies and constraints	Approved help content, a controlled evaluation environment, reliable test cases, trained reviewers and clear escalation routes must be available. Privacy, retention, security, identity, least privilege, supplier terms, cost, latency, monitoring and rollback controls are not yet approved.
Proposed next step	The Head of Customer Operations should lead a controlled internal evaluation with the Support Quality Lead, Privacy, Security, Technology and staff representatives. Review the evidence and unresolved controls before deciding whether to seek approval for a limited live 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.

Internal evaluation is workable, but live-pilot controls are missing

The proposed discovery-stage use case has a clear problem, owner and bounded internal evaluation plan. It is ready to continue design and controlled testing, but several live-pilot governance and control items are still unresolved.

Findings

- The business problem is well framed for discovery work: manual classification, routing and drafting take time and create inconsistency.
- The use case has clear success measures, including first response, re-routing, policy errors and mandatory escalation detection.
- The current plan does not yet name the senior AI-governance owner or the operational incident owner needed for any live pilot.
- Customer correction and challenge pathways are not approved.
- Cost viability cannot yet be judged because no budget, supplier cost or staffing impact is provided.
- The proposed next step is a controlled internal evaluation using synthetic and approved de-identified cases, which matches the current discovery stage.

Evidence used

- Confirmed use-case version 1
- E-03-operating-boundaries-and-oversight.md

Limitations

- No independent operational evidence is available yet for live performance.
- Cost viability cannot be rated from the supplied records.
- This assessment is limited to the business case and discovery-stage fit.
- The supporting evidence is a draft operating boundary, not proof that controls work in practice.

Review reference: REV-11649850

02 People assessment

Evidence strength: Limited

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

Discovery is suitable for a controlled internal evaluation

The use case has a clear people-impact scope, a defined human decision boundary and a named owner for the discovery stage. Key gaps remain on staff and customer consultation, correction pathways and named governance roles before any live pilot.

Findings

- The people affected by the proposal are clearly identified for discovery stage work.
- The decision boundary keeps final human control with support agents, team leads and authorised people.
- Consultation with affected customers and frontline staff is not yet shown.
- The current material does not define accessible formats, assistance channels or alternative ways to interact with the process.
- Training, role changes and workload support for support agents and team leads are not yet set out.
- Customer correction and challenge pathways are still unresolved for any live pilot.
- The internal evaluation uses synthetic and approved de-identified examples, which suits a discovery-stage test.

Evidence used

- Confirmed use-case version 1
- E-03-operating-boundaries-and-oversight.md

Limitations

- This assessment is preliminary and limited to the supplied use case and evidence records.
- It does not assess business value, technical feasibility, privacy, security, legal or policy compliance.
- No live-customer results are available, so this review cannot judge operational performance in practice.

Review reference: REV-EAE23E33

03 Technology assessment

Evidence strength: Limited

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

Internal evaluation is workable, but controls are still open

The proposed discovery-stage evaluation has enough definition to proceed to controlled testing. The main blocker for the next step is missing ownership and unapproved controls for any later live pilot.

Findings

- The discovery-stage boundary is clear enough to start a controlled internal evaluation.
- Future live-pilot readiness is not yet established because key ownership and control items remain open.
- The current data boundary is suitable for controlled testing, but it does not yet define how future live data would be retained or deleted.
- Resilience and supportability are not yet measurable because monitoring, latency and rollback arrangements are still unapproved.
- The baseline measures and success criteria are specific enough to evaluate the test outcomes.

Evidence used

- Confirmed use-case version 1
- E-03-operating-boundaries-and-oversight.md

Limitations

- No supporting evidence was supplied for production or pilot performance.
- The assessment is limited to technical and operational feasibility at the discovery stage.
- Customer adoption, business case, policy and legal issues were out of scope.

Review reference: REV-020609C0

04 Risks and oversight

Evidence strength: Limited

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

Discovery can continue, but key live-pilot controls are missing

The proposed internal evaluation has a clear boundary and named human review, so it is suitable to keep designing and prepare controlled testing. The main blockers are for any later live pilot: governance owners, customer challenge pathways, and privacy, retention and security controls are still open.

Findings

- The proposed discovery-stage evaluation has a bounded task and a clear human review step.
- A later live pilot would expose customers to unresolved challenge and correction gaps.
- Accountability for any live pilot is incomplete because key roles are not yet appointed.
- The draft boundary excludes sending replies, account changes, refunds and other high-impact decisions.
- The remaining risk level for a future pilot is still unclear because privacy, retention, security, supplier, monitoring and rollback controls are not yet approved.

Evidence used

- Confirmed use-case version 1
- E-03-operating-boundaries-and-oversight.md

Limitations

- This result is preliminary for any live pilot because the use case is still at Discovery.
- I did not verify technical enforcement of the draft controls, only the documented boundary.
- I did not assess legal compliance or policy profile selection.

Review reference: REV-52D0D651

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

Discovery stage is suitable, but live-pilot controls are still open

The use case is bounded enough for a controlled internal evaluation. Key live-pilot governance, privacy and customer-challenge controls are still missing.

Review status: No stop

Findings

- The policy profile is relevant to a discovery-stage AI support assistant that is starting with lower-risk use.
- The use case has a clear boundary, human review and a limited internal evaluation plan.
- A senior AI-governance owner and an operational incident owner are not yet appointed for any live pilot.
- Customer correction and challenge pathways are still unresolved for any live pilot.
- Privacy, retention, security, supplier, monitoring and rollback controls are not yet approved.

Overlaps

- Both the policy profile and the confirmed use case point to a discovery-stage internal evaluation with human review.
- The four core assessments agree that the internal evaluation is workable but live-pilot controls are still open.

Conflicts for human review

- The use case says the internal evaluation is suitable to continue, but the evidence also shows live-pilot governance and correction pathways are not yet approved.
- The guidance is voluntary, so there is no conflict with law or mandate in the supplied material.

Limitations

- The review is based only on the supplied records and does not test real-world operation.
- The supporting evidence is a draft control boundary, not proof that controls are technically enforced.
- No live-customer results or approved live-pilot controls were 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-943709BF.

04 / DECISION RECORD

Conditions, gaps and limitations

Conditions

- Some live-pilot governance roles are not yet appointed.
- Customer correction and challenge pathways are not approved yet.
- Privacy, retention, security, monitoring and rollback controls still need approval.
- The use case can affect customers in vulnerable circumstances, so it deserves closer review.

Related evidence gaps

Accountability and human decisions

- Future live-pilot readiness is not yet established because key ownership and control items remain open. A senior AI-governance owner and an operational incident owner are not yet appointed for any live pilot.
- Named senior AI-governance owner and operational incident owner for any live pilot, with pause authority and escalation duties.

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

Testing, monitoring and evidence

- A simple cost and resourcing view for the internal evaluation and any limited live pilot.

Contributing saved outputs: Business assessment

People and use

- Accessible interaction and challenge options for customers who need assistance, alternative formats or non-digital support.

Contributing saved outputs: People assessment

Other evidence gaps

- Training and change support arrangements for support agents and team leads.
- A consistent method for measuring mandatory escalations.

Contributing saved outputs: People assessment; Technology assessment

Data and information

- Approved privacy, retention, deletion, access and least-privilege controls for any live data use.
- Defined monitoring, latency and rollback arrangements for live operation. Approved privacy, retention, security, supplier, monitoring and rollback controls for any live pilot.

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

Limitations

- I did not assess legal compliance or policy profile selection.
- The supporting evidence is a draft operating boundary, not proof that controls work in practice.
- This result is preliminary for any live pilot because the use case is still at Discovery.
- It does not assess business value, technical feasibility, privacy, security, legal or policy compliance.

- 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	116498507e744a8401b09175f639a460affa4b07977d9283773bddae683598b	viewed
People assessment	eae23e33b2b100bb1ec9ada57e0b6b573609ecdcc9ec7a058c52590fc093a064	viewed
Technology assessment	020609c01898d855c3bd41ba4b4476d97f81015fb72f3bc284abfce5812fc9ee	viewed
Risks and oversight	52d0d6515f4b202bb0d65056998afc8b919bb127339afbafbddcc5e15b2b398	viewed
Guidance for AI adoption: foundations	Published 5 May 2026 · 943709bdfdc02351b0c9e48a31d6f3e554ff6161a5022efa137ed18d64abdb76	viewed

Report identity: fa56932d44836f7d74cf00b0fe3c83b1af6bce7e41d5d850070499c23bd2784b

Confirmed use-case version: 1. Source assessment revision: 20.

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