

AI Leasing Assistant: User Stories & Acceptance Criteria

Representative backlog slice defining supported renter conversations, clarification, human escalation, context handoff, and AI boundaries.

Artifact	User stories + acceptance criteria	Product area	AI renter engagement / leasing
Prepared by	Nina Bernardo - Senior Business Analyst	Focus	Behaviour, boundaries, escalation, edge cases

Confidentiality note. Representative reconstruction based on the type of AI/product requirements work performed professionally. Product rules, intents, examples, data and implementation details have been generalized or changed.

1. Product behaviour being defined

The assistant should handle common renter questions when it has enough approved context, ask a useful clarifying question when intent is ambiguous, and hand the conversation to a human when the request is unsupported, sensitive, explicitly asks for a person, or cannot be answered confidently from available information.

Intent / situation	AI behaviour	Human escalation?	Notes
Property / unit information	Answer from approved property/unit context when available.	Not normally	Do not invent unavailable unit details.
Tour / leasing next step	Guide renter through supported next step or collect required information.	Conditional	Escalate if requested action is outside supported workflow.
Ambiguous renter question	Ask one targeted clarification before choosing a path.	Not yet	Avoid guessing between materially different intents.
Unsupported policy / exception	Do not fabricate a policy answer.	Yes	Transfer with conversation context.
Sensitive / high-risk request	Stop autonomous handling according to defined boundary.	Yes	Use approved handoff language.
Renter explicitly asks for a person	Acknowledge and initiate handoff.	Yes	Do not keep trying to retain the conversation.
Repeated failure to understand	Stop looping after defined threshold.	Yes	Preserve what has already been learned.

2. Representative backlog

US-01 — Answer a supported renter question

As a prospective renter, I want to receive a useful answer to a supported property or leasing question so that I can continue my search without waiting for an agent.

AC-01 — Given the renter asks a supported question and the required approved context is available, when the assistant processes the message, then it provides an answer grounded in that available context.

AC-02 — Given the required information is not available, when the renter asks the question, then the assistant does not invent the missing information.

AC-03 — Given the renter asks more than one supported question in the same message, when enough context is available, then the assistant addresses each supported part or clearly identifies what it cannot answer.

US-02 — Clarify ambiguous intent

As a prospective renter, I want the assistant to clarify what I mean when my request could reasonably map to different actions so that I am not taken down the wrong path.

AC-01 — Given the renter message maps to two or more materially different supported intents and context does not resolve the ambiguity, then the assistant asks a targeted clarification before executing either path.

AC-02 — Given prior conversation context already resolves the ambiguity, then the assistant uses that context rather than asking the renter to repeat information.

AC-03 — The clarification should request only the information needed to choose the next path and should not restart the conversation.

US-03 — Escalate when a human is needed

As a prospective renter, I want to be connected to a person when the assistant cannot appropriately handle my request so that I can still get help without starting over.

AC-01 — Given the renter explicitly asks to speak with a person, then the assistant acknowledges the request and initiates the supported human-handoff path.

AC-02 — Given the request falls into a defined unsupported or sensitive category, then the assistant does not attempt to generate a substantive answer before escalation.

AC-03 — When escalation occurs, the handoff includes the available conversation context required to reduce unnecessary repetition by the renter.

AC-04 — If the handoff mechanism is unavailable, the assistant provides the approved fallback rather than claiming that transfer succeeded.

US-04 — Stop repeated misunderstanding

As a prospective renter, I want a useful fallback when the assistant repeatedly fails to understand me so that I am not trapped in a loop.

AC-01 — Given the assistant has asked for clarification up to the agreed threshold without resolving the request, then it stops repeating the same clarification pattern.

AC-02 — At the threshold, the assistant offers the approved human or alternate-support path.

AC-03 — Information already supplied by the renter remains available to the handoff where technically supported.

US-05 — Preserve context across the conversation

As a prospective renter, I want the assistant to remember relevant information I already provided during the active conversation so that I do not have to repeat myself.

AC-01 — Given the renter has already supplied a relevant value during the active conversation, then the assistant uses it when processing a later related request.

AC-02 — The assistant must not treat unrelated prior content as confirmed information for the current intent.

AC-03 — If two pieces of conversation context conflict, the assistant asks for confirmation rather than silently selecting one.

US-06 — Handle unavailable or uncertain information safely

As a prospective renter, I want the assistant to be clear when it does not have reliable information so that I do not make a decision based on a fabricated answer.

AC-01 — Given the assistant lacks approved data required to answer, then it does not present a guessed value as fact.

AC-02 — Where an approved alternate path exists, the assistant offers that path or escalates according to the defined rule.

AC-03 — The assistant does not imply that an unavailable action, appointment, unit, price, policy, or outcome has been confirmed.

3. Edge cases raised during refinement

ID	Edge case	Question / expected direction
EC-01	Renter asks one supported question and one unsupported question in the same message.	Answer the supported portion only if doing so does not create a misleading partial response; escalate/redirect the unsupported portion.
EC-02	Renter changes intent mid-conversation.	Use the latest clear intent; do not force completion of the previous flow.
EC-03	Renter says 'yes' after several possible questions were asked.	Avoid assuming which question was answered if the referent is unclear.
EC-04	Renter asks for a human, then immediately asks another simple question.	Human request remains authoritative unless product explicitly defines cancellation of handoff.
EC-05	AI has partial property context but not enough to confirm the requested detail.	State the limitation or use the approved fallback; do not fill the gap from general knowledge.
EC-06	External action fails after the assistant tells the renter it is being performed.	Do not claim completion; communicate the failure/fallback according to the integration outcome.

4. MVP boundary / deferred decisions

MVP / include	Defer or resolve separately
Supported high-volume renter intents with approved source data	Broad open-domain answering
Targeted clarification for ambiguous supported intents	Complex multi-step reasoning not tied to a defined renter journey
Explicit human request + unsupported/sensitive escalation	Autonomous handling of policy exceptions requiring human judgment
Context handoff and loop prevention	Assuming every integration/action can recover automatically
Failure behaviour for supported actions	Expansion into additional channels/intents before core paths are validated

What this sample demonstrates: AI requirements need more than a list of intents. The delivery detail is in defining when the assistant can answer, when it needs more information, what context it can rely on, what it must never assume, and when it should stop trying and hand off.