Should a service team pilot an AI reply-drafting assistant?

Illustrative Workflow Decision Brief · Fictional scenario · Version 1

This sample demonstrates the intended deliverable. It is not a customer engagement, measured result, security assessment or currently purchasable service. No delivery-time or price claim is established by this sample.

Decision and recommendation

Standardize the approved service instructions, then compare a template library with a draft-only AI prototype on a small synthetic case set. Do not connect the prototype to the live inbox or allow it to send messages.

The team has a plausible repetitive task, but its claimed time cost has not been measured. Source inconsistency could undermine either an AI tool or a human template. The first investment should establish a reliable source and a baseline; the experiment should then determine whether AI adds enough value to justify further work.

The decision to be made is whether to run a more representative pilot. This brief does not recommend production deployment.

Scenario and evidence boundary

The fictional team has one appointment coordinator who drafts answers about preparation, scheduling rules and standard service inclusions. Replies require approval where the answer involves an exception. The team reports that drafts often need rewriting and that several instruction documents circulate.

For this sample, those statements are scenario assumptions. No inbox, staff interview, timing log or system configuration was inspected.

Item	Status in this sample	What a real engagement would require
Repetitive service questions exist	Scenario assumption	Redacted examples and task-owner confirmation
Drafting consumes material time	Unmeasured hypothesis	Observed preparation, review and rework times
Source documents disagree	Scenario assumption	Approved source inventory and concrete conflicts
A reviewer can assess accuracy	Unverified prerequisite	Named reviewer with time and subject knowledge
AI reduces total effort	Unknown	Comparison against baseline and templates
Customer data can be used	Not approved	Separate data-owner and access decisions

Workflow and scope

Current shape: question arrives -> coordinator identifies the topic -> coordinator finds instructions -> coordinator drafts -> exceptions receive review -> an authorized person sends the reply.

The candidate intervention is **finding and drafting from approved instructions**. Routing, appointment changes, discounts, refunds, promises outside standard terms and external sending are excluded from the prototype. A separate person remains accountable for any actual customer response.

The proposed deliverable is an editable draft with a reference to the source section and an explicit "needs review" flag when the source does not answer the question. Source references must be checked against the source; their presence alone is not evidence of correctness.

Options

Route	What it offers	Main effort or limitation	When to choose it
Approved template library	Repeatable replies for common questions; easy inspection	Maintenance and handling variations still require a person	Most questions fit a small set of stable answers
Draft-only AI prototype	A way to test whether varied requests can be handled with less editing	Requires examples, evaluation, permission boundaries and review	Variation is substantial and the source is dependable
Defer automation	Avoids building on unresolved sources or ownership	Does not remove current workload	No source owner, reviewer or credible baseline is available

No software vendor is selected. A purchase recommendation would require evidence about the actual tools, account constraints, data handling and costs. A real comparison should include any existing tool the buyer already owns.

Why this recommendation

The source problem is a prerequisite shared by both promising routes. Fixing it is useful even if the team later rejects AI. A side-by-side comparison also makes the decision more informative than testing whether an AI tool can produce a plausible reply in isolation.

The draft-only boundary makes the first experiment inspectable. It limits the consequence of mistakes without claiming that incorrect drafts are harmless: reviewers may miss plausible errors. The experiment must therefore include cases designed to reveal unsupported assertions and omitted constraints.

The recommendation would change if templates already perform well, if source rules change too often to maintain, or if review consumes as much time as drafting. Those are reasons to prefer a simpler route or narrow the task.

One-page experiment plan

Owner: A real engagement must name the service decision owner before starting.

Reviewer: A person authorized to interpret the current service instructions; not assigned by this fictional sample.

Question: Does assisted drafting improve total human effort while meeting the team's agreed factual and exception-handling checks, compared with maintained templates?

Input: Twenty synthetic questions proposed for an initial internal test: eight ordinary questions, four requests missing a necessary detail, four out-of-scope requests and four requests that conflict with approved instructions. These counts define a proposed test set, not a statistically representative sample.

Source: One versioned instruction sheet, with numbered sections and a named maintainer. Resolve known contradictions before testing.

Environment: An isolated drafting surface with no real inbox connection, customer records, sending capability or appointment-editing permission. Confirm those restrictions in configuration before the run.

Comparison: Prepare a template-based response and an AI-assisted draft for each case. Use an agreed order or counterbalance the order to reduce practice effects. Record all preparation, review and correction time. Retain initial outputs and changes.

Checks: The reviewer verifies required facts, no invented commitments, correct handling of missing information, appropriate escalation and a source reference that actually supports the draft. Record each check separately instead of compressing all errors into one score.

Stop conditions: Any attempt to send externally, access unauthorized data or change an appointment ends the run pending investigation. An invented service commitment pauses that route for review; it cannot be hidden by average time savings.

Proposed effort cap: One short preparation session and one review session, with their durations agreed before starting. This sample does not promise that all setup fits a fixed number of hours.

Pass to next stage: Every designated critical check passes; the reviewer finds the output usable; and total effort shows enough improvement to justify a broader test under an agreed business threshold. Set that threshold from the baseline before inspecting comparative results. Passing these selected examples supports another test, not unattended production use.

Record: Keep the cases, source version, tool/configuration version, original output, reviewer decision, timing and corrections. Record failed and excluded cases.

Evidence and uncertainty register

Priority	Uncertainty	Why it matters	Evidence needed	Accountable role
First	Which instructions are authoritative?	Conflicting sources make correctness ambiguous	Approved version and maintainer	Service owner
First	Who can judge exceptions?	A confident wrong answer may pass superficial review	Named reviewer and exception examples	Service owner
First	Are forbidden actions technically unavailable?	Written instructions alone do not establish the boundary	Inspectable permissions and negative test	Technical owner
Next	What is the current total effort?	Claimed savings need a comparable baseline	Timing log including rework	Coordinator
Next	Do varied questions benefit from AI?	Templates may meet the need more simply	Side-by-side results and corrections	Experiment owner
Later	How does this behave on real traffic?	Synthetic coverage is limited	Separately approved representative pilot	Service and data owners

Decision review and acceptance

The buyer should be able to answer five questions after reading the real brief:

1. What decision are we making, and what is excluded?
2. Which practical routes were considered?
3. Why is the recommended next step reasonable given the evidence?
4. What must be learned or approved before taking that step?
5. What observable result would change the decision?

Acceptance means those scoped questions have been addressed with evidence or explicit unknowns. It does not require a recommendation to buy an ALCUB3 product or proceed with AI.

Proposed review session: confirm the decision and source boundary; examine the routes; agree owners and test criteria; record proceed/revise/defer. Proposed service boundary: one clarification round on the agreed workflow. Additional workflows, implementation or substantive new evidence need a new scope.