Worked Example: HR Policy Answer Drafting

Follow one fictional workflow from opportunity to a pilot and operating decision.

For: HR operations teams, product owners, and sponsors | Use when: Alongside the Five Decisions and value worksheets | Suggested time: 15 minutes

Version 1.0 | October 5, 2026

All organizations, workflows, numbers, and results in these examples are fictional and provided for learning. Targets are proposed decision criteria, not validated standards or promised outcomes.

The situation

A fictional HR service desk handles 300 routine leave-policy questions per month. Analysts search different document versions and send answers after specialist review. The proposed product drafts a cited explanation; people continue to approve and send it.

Completed Five Decisions

The smallest useful result is a supported draft or a clear escalation, not an autonomous HR decision.

Problem

Assume 15 minutes active handling per question, including review. Repeated lookup and policy-version confusion are the first hypotheses to validate.

Product

Draft with current policy citations. Exclude personal eligibility, personnel records, automatic sending, and policy interpretation beyond the approved source.

People

Analyst uses; policy specialist reviews; borrowed platform specialist prototypes; service-desk lead owns the workflow and sponsor owns expansion.

Production

Approved synthetic corpus first. One policy owner maintains versions. Missing or conflicting sources route to the specialist. Manual lookup remains the fallback.

Proof

Run 40 synthetic cases: 20 routine, 5 incomplete, 5 conflicting-source, 5 individual-eligibility, 5 unavailable-source. Investigate any unsupported claim; block expansion for missed eligibility escalation or unauthorized action.

Completed route and responsibility decisions

Choose Borrow for a scoped retrieval prototype. Compare Rent before approving a custom production system. The service-desk lead owns ongoing support; the platform partner owns technical maintenance. AI may retrieve and draft. It may not decide eligibility or send. The specialist checks evidence and approves every pilot draft.

Completed value and cost worksheet

These are planning assumptions, not measured results. Use identical case mix and total active handling time in both baseline and pilot.

Efficiency target

300 questions/month; proposed change from 15 to 9 active minutes per question.

Capacity

300 x (15 - 9) / 60 = 30 hours/month released if the target is achieved.

Financial mechanism

No spending reduction is established. The 30 hours are capacity, not cash savings.

Cost horizon

Illustrative one-time cost $3,000; incremental recurring cost $250/month. First three months cost $3,000 + (3 x $250) = $3,750. Count review effort in handling time; do not add it twice.

Quality and risk

Review citation support, corrected claims, and escalation behavior. Faster answers cannot compensate for unauthorized decisions.

Illustrative test result and decision

Suppose the test produces supported behavior in 19 of 20 routine cases and all 20 exception cases. One routine draft cites a stale policy. Even though 39 of 40 cases meet expectations, the owner chooses REVISE: remove the stale source, add version checks, and re-test affected and neighboring cases. The run does not establish readiness for real employee questions.

Completed 30/60/90 plan

Days 1-30: validate the real baseline through an authorized process, clean sources, test synthetically, and resolve blocking defects. Days 31-60: only if approved, run a limited pilot with review of every draft. Days 61-90: inspect quality, effort, adoption, costs, and readiness; expand, revise, or stop. No commitment to autonomous sending or eligibility decisions.

Original guide: https://builtinto.work/resources/worked-example-policy-workflow