



SAMPLE CLIENT DELIVERABLE

Customer Service Copilot Implementation Plan

A six-week, human-in-the-loop pilot for a retail and e-commerce operation.

PREPARED BY

Keith Soyka

AI Implementation & Systems Advisory

PLAN HORIZON

6-week pilot

with controlled expansion gates

Build the smallest system that can earn its keep: faster, better-grounded answers for support agents, with people retaining the final decision and customer relationship.

Illustrative scenario. This sample demonstrates the structure and depth of a client implementation plan. The company, workflows, volumes, targets, and economics are modeled assumptions and do not represent a past client's results.

RECOMMENDATION

Start narrow, measure honestly, and keep the human in control

The proposed pilot equips service agents with grounded draft answers drawn from approved policy, product, and order sources. It does not send messages, issue refunds, or make exceptions without an agent's review.

6 weeks

PILOT DURATION

Discovery through handoff, with a decision gate every week.

4

INITIAL INTENTS

Order status, returns, product questions, and escalation support.

100%

HUMAN REVIEW

Every customer-facing response remains agent-approved during the pilot.

What this plan solves

- ✓ Agents hunt across storefront, help desk, policy files, and order records.
- ✓ Routine questions absorb time needed for exceptions and relationship recovery.
- ✓ Answer quality varies with tenure and source familiarity.
- ✓ Leaders lack a controlled way to test automation without risking trust.

What success looks like

- ✓ Drafts cite or link to the approved source used.
- ✓ Unsupported answers and policy exceptions route to people.
- ✓ Median handling time improves without lowering resolution quality.
- ✓ Agents understand when to accept, edit, or reject a suggestion.

Decision rule: expand only if quality, safety, speed, and agent adoption all clear the agreed threshold. A faster system that creates rework or weakens customer trust does not pass.

Working assumptions to confirm in discovery

The retailer uses a commerce platform, a ticketing/help-desk system, documented return and shipping policies, and identifiable subject-matter owners. Volumes, baseline performance, integration access, and legal requirements remain unverified until kickoff.

PILOT BOUNDARY

The first release assists decisions; it does not take them over


In scope

- Retrieve approved policy and product content.
- Summarize relevant order context.
- Draft answers for four initial intents.
- Suggest an escalation reason.
- Capture accept, edit, reject, and source-use events.


Explicitly out of scope

- Autonomous customer messaging.
- Refunds, credits, cancellations, or price changes.
- Policy overrides and fraud decisions.
- Unapproved web search or staff documents.
- Expansion to marketing or merchandising workflows.

Target workflow



Reference architecture



Data boundary: the model receives only the minimum order and policy context needed for the active ticket. Payment data, authentication credentials, and unrelated customer history are excluded.

SIX-WEEK SEQUENCE

Each week produces evidence and a decision, not just activity

W1	W2	W3	W4	W5	W6
Discover + baseline	Govern sources	Build workflow	Evaluate safely	Controlled pilot	Handoff + decide

WEEK 1	Confirm the operating truth Map the current ticket path; sample intents; establish baseline handling time, resolution quality, escalation rate, and rework; confirm access and privacy boundaries.	Gate: scope + baseline approved
WEEK 2	Prepare trusted knowledge Inventory sources, assign owners, resolve contradictions, tag effective dates, define freshness rules, and create a reference answer set.	Gate: source set accepted
WEEK 3	Configure the assisted workflow Connect minimum required systems; implement retrieval, drafting, source display, escalation rules, event logging, and role-based access.	Gate: test environment ready
WEEK 4	Test before real use Run the reference set plus edge cases; review unsupported claims, policy conflicts, privacy exposure, escalation accuracy, latency, and usability.	Gate: safety + quality pass
WEEK 5	Pilot with a small agent cohort Train selected agents; monitor daily; collect accept/edit/reject reasons; compare against baseline while maintaining full human review.	Gate: pilot evidence sufficient
WEEK 6	Stabilize, document, and decide Resolve priority defects; finalize operating guide and administrator runbook; report outcomes; recommend scale, revise, or stop.	Gate: owner signs handoff

Schedule dependency: this sequence assumes timely source-owner decisions and access to a test environment. Delayed inputs move the gate; they do not justify bypassing it.

OPERATING MODEL

Implementation succeeds when ownership is visible

WORKSTREAM	ADVISOR LEADS	CLIENT OWNS	EVIDENCE PRODUCED
Workflow	Process map, pilot design, exception logic	Operational truth, priority intents, policy decisions	Approved current/future-state maps
Knowledge	Source inventory, preparation method, retrieval tests	Source authority, corrections, freshness owners	Governed source register + answer set
Technology	Configuration, integrations, logging, test harness	Access approvals, vendor contacts, environments	Working pilot + technical runbook
Risk	Control design, failure tests, incident workflow	Privacy/legal interpretation, risk acceptance	Control checklist + issue log
Adoption	Training, feedback loops, usage analysis	Agent participation, coaching, operating changes	Training pack + pilot feedback
Measurement	Baseline method, dashboard definition, readout	Metric approval, business interpretation	Scorecard + decision memo

Core project roles

- **Executive sponsor:** resolves scope and risk decisions.
- **Support owner:** owns workflow and pilot-agent readiness.
- **Knowledge owner:** approves source authority and updates.
- **Technical owner:** grants access and accepts the runbook.
- **Advisor:** designs, configures, tests, measures, and transfers.

Decision cadence

Twice weekly: delivery working session and risk review.

Daily in pilot: issue triage, with severity-based escalation.

Every gate: named approver records proceed, revise, or stop.

No hidden dependency: the client receives editable documentation, source ownership, configuration notes, test cases, and operating instructions. Continued advisory support is optional, not required to keep the pilot understandable.

Change control

New intents, systems, decision rights, or autonomous actions are treated as scope changes. They enter the backlog with impact, risk, effort, and owner documented before approval.

ACCEPTANCE FRAMEWORK

The pilot must prove quality, safety, speed, and adoption together

DIMENSION	MEASURE	ILLUSTRATIVE PILOT THRESHOLD*	HOW IT IS CHECKED
Quality	Grounded answer accuracy	≥ 90%	Human-scored reference set and sampled pilot drafts
Safety	Critical policy/privacy failures	0	Edge-case tests plus incident review
Escalation	Required escalations correctly flagged	≥ 95%	Labeled test cases and ticket sampling
Speed	Median handling-time change	≥ 15% better	Comparable pilot vs. baseline intents
Adoption	Useful draft rate	≥ 70%	Accepted or lightly edited drafts / eligible drafts
Trust	Agent confidence	≥ 4 / 5	Structured pulse survey plus interviews

*Targets are examples for this fictional scenario. Actual thresholds require a measured baseline, risk tolerance, and sponsor approval.

Required controls before live pilot

● Wrong policy	Approved-source allowlist, effective dates, citations, refusal on weak evidence	Knowledge owner reviews conflicts and signs the source register
● Privacy exposure	Minimum-data retrieval, field restrictions, role-based access, retention review	Test with synthetic records before authorized production data
● Over-reliance	Visible evidence, mandatory agent action, confidence-free interface language	Training emphasizes accept, edit, reject, and escalate behavior
● Silent drift	Versioned prompts and sources, change log, regression set, usage monitoring	Re-run evaluation after material policy or workflow changes
● Service failure	Manual fallback, health checks, incident owner, fail-closed behavior	Agents retain the existing workflow throughout the pilot

Stop condition: pause the pilot immediately for a critical privacy breach, unauthorized action, repeated unsupported claims, or a failure that prevents agents from verifying the source.

EXIT WITH A MAINTAINABLE SYSTEM

The handoff package makes the next decision legible

Deliverables

- ✓ Validated pilot configuration and integration map
- ✓ Governed source register with named content owners
- ✓ Reference answer and regression test set
- ✓ Agent operating guide and escalation matrix
- ✓ Administrator runbook and change log
- ✓ Training session and recorded questions
- ✓ Pilot scorecard and evidence-backed recommendation

Acceptance at handoff

- ✓ Named owner can update approved knowledge.
- ✓ Named owner can review logs and pilot metrics.
- ✓ Agents can use the manual fallback.
- ✓ Open defects carry severity, owner, and disposition.
- ✓ Access is reviewed and temporary privileges removed.
- ✓ Sponsor signs the final decision and risk posture.

Final decision framework

SCALE

All gates pass

Expand one intent or team at a time, preserving the same evaluation and control pattern.

REVISE

Value appears, a gate misses

Fix the narrow failure, rerun the affected tests, and extend only with explicit approval.

STOP

Trust or economics fail

Preserve findings, remove access, and do not force a technology into a workflow it cannot serve.

Inputs required to convert this sample into a client-ready plan

Operations: ticket volumes by intent, baseline service metrics, staffing model, peak periods, escalation paths.

Systems: commerce and help-desk platforms, integration constraints, environments, identity and access model.

Knowledge: policy authority, product content, update owners, contradiction and expiration history.

Risk: privacy requirements, data residency, retention, vendor review, approved use boundaries.

Economics: loaded support cost, implementation budget, expected maintenance capacity, value threshold.

People: sponsor, support owner, technical owner, pilot agents, legal/privacy reviewer.

Recommended next action: run a focused discovery and baseline sprint, then issue the final scope, acceptance thresholds, schedule, and fixed implementation estimate from verified inputs.