


## FIELD BRIEFING

# Give the pilot somewhere to land

Give the pilot somewhere to land.



## What turns an interesting experiment into a maintained workflow?

A pilot should produce a decision, not an indefinite demonstration. Name the workflow, its users, its owner and the result you want to learn. Keep the experiment small enough that inputs, outputs and failures can be inspected.

Explanatory schematics and worked examples are illustrative. No live model, observed client results or measured market performance are presented.

# Read the mechanism.

An invented service team wants help preparing enquiry replies. The pilot uses synthetic and approved examples, creates drafts only and records review time. Its success criteria cover correct policy use, no invented promises, useful tone and the effort required to check the draft. Sending remains a separate permission.

First establish the current process. Who prepares a reply, which documents they consult and how long checking takes are part of the baseline. Then test the proposed workflow on representative cases. Finally, review the comparison with the people who will own its day-to-day operation.

Use gates rather than a theatrical countdown. The preparation gate confirms information rights and a reviewer. The evidence gate reviews results and failure types. The adoption gate assigns ownership, support and a rollback plan. The same sequence may take days or months; a calendar does not guarantee readiness.

Document what changed. A new model, revised policy, prompt edit or extra tool may affect behaviour. Keep the versioned configuration and case set so you can repeat the evaluation. Record known limitations where users will see them, rather than leaving them in a development notebook.

Plan recovery before widening use. If the tool is unavailable, can the existing manual process continue? If a draft contains an unsupported commitment, who catches it? If data was incorrectly included, who handles the incident? A workflow is more credible when those questions have owners.

Retire a pilot that does not justify its overhead. Improving the input documents, using a template or narrowing the task may be a better next step than adding autonomy. Success can mean learning that a particular application is not worth adopting.

Technical references: Microsoft · Responsible AI; Anthropic · Building Evals; OpenAI Agents SDK · Human-in-the-loop. Links and source scope appear in the final section.

---

# Try the idea.

## INVENTED TEACHING EXAMPLE

### An enquiry-draft pilot ends in a decision

Prepare: name the task, source rights, reviewer and baseline. Test: review representative drafts for unsupported promises, omissions and checking effort. Adopt or stop: assign support and rollback, then decide whether the observed benefit justifies the overhead. A deadline is not evidence that a gate has passed.

# Go deeper.

## A gate has evidence and an owner

A gate is a decision point, not just a date on a roadmap. Preparation requires input rights and a reviewer; adoption requires observed results, support and rollback ownership.

## Keep the tested system identifiable

Record model, prompt, policy version, data scope and tool configuration. Without that record, a result can be hard to reproduce or interpret after an update.

## Stopping is a result

A pilot can justify narrowing a task, switching to ordinary software or not adopting the workflow. Record why: review overhead, inadequate evidence, error consequence or lack of useful benefit.

# Keep the evidence close.

---

## Pilot charter

A short record of the task, owner, scope and decision the experiment should inform.

---

## Adoption gate

A point at which evidence and ownership are checked before widening use.

---

## Rollback

A planned return to a workable earlier process.

---

## Re-evaluation trigger

A meaningful change that requires the tested assumptions to be checked again.

---

## Sources · accessed 4 October 2026

---

[Microsoft · Responsible AI](#)

---

[Anthropic · Building Evals](#)

---

[OpenAI Agents SDK · Human-in-the-loop](#)

Technical and security references support the mechanisms, not our fictional results. KAF's worksheets and examples are original editorial recommendations. No current provider pricing, comparative ranking or legal compliance is asserted. Verify the exact release, configuration and policy for your use.

---

Practical AI Insights · First editorial edition · 4 October 2026

Give the pilot somewhere to land · [Web edition](#)