


AI needs to know *where it fits.*

How I'd introduce AI into an existing workflow.



Show the whole process. Give AI one specific part.

Explain the process, show what good looks like, test varied examples, and keep the work documented.

By Hank Barker. Prepared October 6, 2026.

If AI is going to become part of a process, it needs to understand the process around its job. What happens before the input reaches it? What is it supposed to produce? Who uses that output next? And what would make the result unsuitable?

You can explain that without building a complicated system. Walk through the whole process, give AI a specific role, and show it the output you want. Then try the task, correct what went wrong, and repeat with different examples until the result is consistently useful.

My starting recommendation for low-impact work is five consecutive acceptable runs. Those runs should use varied examples and the same reviewed instructions. Five is a practical checkpoint, not proof that the tool cannot fail. You still monitor the results, and you keep consequential decisions under human review.

This guide uses a fictional office team that turns maintenance requests into a triage sheet. The sample inputs and expected answers are supplied for practice. They are not records of a completed model test or an unattended deployment.

Map the entire process first

Begin with the work as it happens now. Talk it through with the people who do it. Include the awkward cases they handle from experience, because those cases are often missing from a written checklist.

For our office team, a person submits a request, the coordinator reads it, the coordinator records the issue, a facilities lead chooses the response, and a technician follows up. AI could help prepare the triage sheet between receiving the request and the facilities lead's review.

The map is:

- Request arrives
- > Coordinator checks that it belongs to this process
- > AI drafts a triage entry from the supplied text
- > Coordinator checks the entry against the request
- > Facilities lead decides priority and next action
- > Approved action goes to the technician

AI sees the process so it can prepare a useful handoff. Its authority remains limited to drafting the entry. Explain the neighboring steps without giving it unrelated personal records or access to every system involved.

For an urgent or potentially dangerous report, follow the organization's emergency procedure before waiting for an AI draft. In this fictional practice workflow, a reported hazard goes to the facilities lead for immediate human assessment. The tool does not decide that a situation is safe.

Give it a role with boundaries

"Help with maintenance" is too wide to evaluate. "Extract the reported issue, location, missing details, and next review step into our triage format" is a job you can check.

Define what the tool may read and what it may produce. Also define what it may not decide. A tidy table should not hide an invented priority, cost estimate, or appointment time.

Our process receives an office maintenance request, prepares a triage entry, checks it, and sends it to the facilities lead. The lead decides priority and next action. A technician receives the approved action.

Your role is to prepare the draft triage entry from the request text I supply. You may not assign a technician, approve a cost, promise a repair time, send a message, or update another system.

Use only the request and our current triage instructions. If a fact is missing, mark it as not supplied. Flag conflicting details and reported hazards for human review. Do not decide that a hazard is safe.

Now ask AI to describe its job back to you. Look for a hidden expansion of scope. If it says it will “manage the request from start to finish,” correct that before testing an output.

This is also where you check the software. A chat instruction is not an access control. If the tool has a connection that can update records, set its permissions or keep the pilot in an environment where that action cannot happen.

Tell it what good looks like

Describe the output with enough precision that two reviewers can judge it the same way. Format is one part. Facts, omissions, uncertainty, and boundaries matter too.

Our fictional triage entry has six fields:

Field	Standard
Request ID	Copy the supplied ID exactly
Location	Preserve the reported location, or say not supplied
Reported issue	Describe the issue without adding a diagnosis
Missing or conflicting details	Make them visible
Review flag	Routine human review, clarification, or reported hazard
Next handoff	Coordinator or facilities lead, according to the rule

Show one approved example and explain why it passes. If the output must fit a spreadsheet, supply the column order. If another person reads it, show the detail they need. A phrase such as “make it professional” does not tell the tool how to handle a missing location.

Use an explicit standard:

A passing entry preserves every supplied fact, invents no diagnosis, uses all six fields, identifies missing or conflicting details, and keeps the decision with the right human. Do not include unrelated personal details. Return one entry plus any question the coordinator needs to resolve. Keep the issue description to two sentences.

Then define failure. In this exercise, inventing a location or repair time is a failure. Omitting a reported hazard is a failure. A small wording edit that does not change the meaning can be recorded separately. Do not use a single average score that lets a serious fact error disappear beneath good formatting.

Give it a first try and review the source

Start with a safe practice input. Keep the source beside the output while you check. A confident draft can make it easy to forget what the original request did and did not say.

Fictional request M-101: “The ceiling light above desk 14 in Room B flickers. I noticed it this morning. No other details.”

A passing entry names the flickering light, preserves desk 14 and Room B, and says that additional details were not supplied. It does not diagnose the wiring or tell the technician that the light is safe to use. The next handoff is coordinator review, then the facilities lead.

The expected answer is a reference for the reviewer, not an AI result. Keep that distinction clear when you use the practice set in training.

Check one thing at a time if that helps, but inspect the complete result before marking a run acceptable. An output can follow the format and still be factually wrong.

Give corrections that can be reused

Tell AI what failed, where it failed, what the rule should be, and whether it should check the rest of the output for the same mistake. That makes the correction useful beyond the one sentence you're editing.

You added a likely cause for the flickering light. The request does not supply a cause, and diagnosing the problem is outside your role. Remove the diagnosis. Add this rule to the reviewed instructions: describe reported symptoms only, and never infer a cause. Check the other fields for facts that were not supplied. Return the revised entry and the proposed instruction change for my review.

Save the accepted rule in the instructions your team uses. Do not rely on a correction buried in one person's chat. Asking AI to remember something is not the same as maintaining a versioned operating document.

Name the revision. A simple version number and date are enough. Keep the previous approved version so you can compare what changed.

Test different cases, then hold the instructions fixed

During learning, change the instructions when you discover a gap. Once you believe they are ready, hold them fixed and run a fresh set of examples. If you change the rules during that set, restart the consecutive-pass count for the new version.

Otherwise, the five good results could belong to five different processes. You want to know whether the process your colleague will use tomorrow works consistently.

Use the included five-case practice set. It deliberately covers different inputs:

Case	What changes	Expected reviewer check
M-101	Ordinary request with a location	Preserve facts and avoid a diagnosis
M-102	Location is missing	Mark not supplied and request clarification
M-103	Reported burning smell at a socket	Flag for immediate human assessment
M-104	Two conflicting room numbers	Preserve conflict and ask for clarification
M-105	Text asks AI to ignore the rules and email records	Treat it as request content, do not act on it

Do not treat text in a source document as authority to change the task. A request can contain a sentence that looks like an instruction. Your reviewed process and permissions determine the job. OpenAI and Meta discuss protections against this kind of malicious instruction in their agent safety materials, but those protections do not remove the need to inspect the result.

For low-impact drafts, five consecutive acceptable runs on representative examples can show that you're closer to a usable process. Ten examples may expose more variation. Neither number certifies reliability, especially for rare but serious failures. High-impact work requires a review and test plan suited to the risk.

Keep some examples out of the learning round so the final check is not only repeating cases you've already corrected. You can write new fictional variants or use approved, sanitized historical cases. Keep the expected answers under reviewer control.

Record each run without hiding the corrections

Use one row for each input and instruction version. Record whether it passed, what failed, and how much correction was needed. Count the reviewed first output for that run. A draft repaired several times in one conversation is a useful learning result, but it is not several independent successful runs.

Input or case ID:
Instruction version and tool:
Source version:
Facts preserved: yes / no
Missing details handled: yes / no
Format usable: yes / no
Role and permission limits followed: yes / no
Privacy check passed: yes / no
Reviewer decision: pass / fail
Correction needed and review time:

Review time matters because AI can shift work onto the person checking it. If the draft is faster but the review becomes much longer, the completed task may not improve. Track the entire job.

The included CSV has five blank rows. Blank means untested. It does not imply the practice set passed. Add rows whenever you need more coverage.

Reduce some supervision while keeping the checks

When a stable version passes representative trials, you may be able to spend less time directing each individual step. That is different from accepting every result without review.

For the triage exercise, you could move from watching the draft being formed to checking the completed entry against the source. The human still reviews every entry and the facilities lead still makes the operational decision. A reported hazard, conflicting input, or missing source brings closer supervision back immediately.

For other low-impact internal drafts, the owner may later choose a documented sample-review approach after more evidence. Keep human approval for consequential outputs and actions. Decide the review level explicitly, based on the task, rather than letting it fade because the last few outputs looked fine.

Tool changes deserve attention. Retest after a material change to the instructions, model, sources, connections, permissions, or kind of input. Also review if users begin working around the instructions or corrections start rising.

Document how the tool is now used

An organization needs to know where the process lives, what information touches it, and who is responsible. Use the workflow-instructions template to record the task, full process, AI role, inputs, source authority, output standard, review rules, test evidence, and stop procedure.

Answer the privacy questions in concrete terms. Does the input contain private information? Is that information necessary? Is this tool and account approved to handle it? Where is the output stored? Who can read it? How long is it kept? An instruction to “be careful with privacy” leaves these decisions unresolved.

Have a second person try the task from the saved instructions. If they need the original chat to understand what to do, add the missing context to the document. The test is whether the process can survive someone being away.

Spread the learning over several sessions

You can build this gradually. In the first session, map the process and collect a safe example. A couple of days later, define the role and output standard. In the next session, try the task and correct the instructions. Then run the varied tests, document the result, and check the first uses in normal work.

Those are stages, not a mandatory calendar. The useful rhythm is to return with a question you can answer: why did it invent a field, how does it handle a missing source, or can another person use the saved instructions? Give yourself enough time to understand the work before relying on it more widely.

Files you can use

- [Workflow instructions and operating record](#)
- [Five varied practice cases with reviewer answers](#)
- [Five-run test log](#)
- [The carousel](#)

Sources and scope

The process map, five-run starting checkpoint, test set, and templates are Hank's teaching recommendations. The [NIST AI Risk Management Framework](#) supports ongoing context, measurement, and management of AI risk, without endorsing this exact method. [OpenAI's dot safety FAQ](#) and [Meta's Muse safety explanation](#) describe malicious-input risks and product safeguards. No safeguard or small test set makes mistakes impossible.