

Making Room for AI

FROM IMPLEMENTATION
TO TRANSFORMATION



JOSHUA BLOODWORTH



Making Room for AI

From Implementation to Transformation

Joshua Bloodworth

A practical field guide for putting AI to work, helping people use it well, and deciding what should change next.

A path through the book

Introduction · Where does AI belong? 4

Part I · Implementation 7

 Chapter 1 · Choose a Useful Starting Point 8

 Chapter 2 · Give AI an Assignment 12

 Chapter 3 · Build the First Working Example 18

Part II · Enablement 24

 Chapter 4 · Decide What Can Change 25

 Chapter 5 · Help People Work Well With AI 29

 Chapter 6 · Connect the Responsibilities 33

 Chapter 7 · Build a Practice That Improves 48

Part III · Transformation 58

 Chapter 8 · Choose What the New Capacity Is For 59

 Chapter 9 · Rethink Roles and Decisions 66

 Chapter 10 · Create New Possibilities 72

Closing · Bring it to the team 79

Workbook · Implementation, Enablement, and Transformation 80

Application index · What AI contributes 111

Working terms 113

Acknowledgments and influences 115

Sources and further reading 116

Where does AI belong?

Maya has a customer call on Thursday. Before it, she needs to find the commitments from the last meeting, check their status with a colleague, and decide what deserves attention now. The information exists, but it sits in notes and an action list. On a busy day, gathering it leaves less time to think about the conversation she wants to have.

An AI-prepared brief could help Maya begin with the relevant records in front of her. That is a useful implementation. The larger question begins when the result works: will her organization give that contribution a dependable place, remove the preparation it replaces, and make room for the better conversation it supports?

A manager can invite people to use AI while leaving every report, meeting, and deadline exactly as it was. People then have to fit learning and experimentation around an already full day, with little authority to change the routine they are supposed to improve. Giving the team room to adjust its work makes that invitation more credible.

That is what I mean by making room for AI. It involves deciding where the capability belongs, helping people work well with it, and being willing to change the arrangements around them. Sometimes a clearer handoff is enough. A more ambitious use may call for a different allocation of responsibility or a service the organization has not offered before. The extent of the change should follow what people are trying to make possible.

The prospect that interests me most is giving people more opportunity to use their judgment and experience. Someone who spends less effort assembling background may have more attention for the person across the table. A small team may be able to explore an idea it has repeatedly postponed. Those are choices to pursue and examine, with the people doing the work helping decide where the additional support would matter.

This book is for people who lead teams, whether they own a small business, manage a department, or coordinate colleagues with more responsibilities than time. It follows one recurring AI assignment from a first test into a team practice and then into a new customer service. You can begin with one activity and decide, with the people involved, how much change the contribution deserves.

Three parts, three different questions

Implementation, enablement, and transformation are used in different ways. Here is how I'll use them in this book.

Part	Working meaning	Question we will explore
Implementation	Putting an AI capability into use for a defined purpose.	How do we put AI to useful work?
Enablement	Helping people develop the skills, judgment, access, and support to use that capability well.	What do people and the organization need to make this work consistently?
Transformation	Making meaningful changes to responsibilities, processes, or services because of what the capability allows.	What can we now do or organize differently?

The sections give us a way to examine different decisions, although those decisions will often overlap. Learning begins with the first example and continues as responsibilities change; permission for an initial trial also needs reconsideration when other people begin relying on it. A team can be redesigning a service while still experimenting with one of the assignments that supports it.

Some assignments in this book begin with a person asking AI for help. Others could be carried out by an agent. I use that term for software set up to use AI and available tools to carry out steps within a defined assignment. What it can actually do depends on its configuration and permissions.

That distinction matters once the work continues beyond one response. A connected assignment may retrieve account records, notice that a status

is missing, ask a colleague one specific question, save its place, and resume when the answer arrives. It may also stop because the answer never comes or because a check finds that the result is not ready. The useful question is not whether the software looks autonomous. It is whether the team can explain what changes the next step, what the system is allowed to do, what sends work to a person, and what counts as completion.

Chapter 6 follows one such assignment through a successful run and an interrupted one. The example is still fictional teaching material. Its purpose is to make the movement of work inspectable before the book asks what that capability could change in the organization.

How to use this book

Parts I and II follow a customer activity from an initial example into a team practice. Our fictional service business includes Maya, who manages customer relationships, Eli, who coordinates service delivery, and Lena, who leads the team. Their customers depend on them for ongoing support and training, and the team wants to make its check-ins more useful.

Part III follows the same team as it develops an earlier planning service for customers, alongside examples from an internal learning team and a group preparing organizational recommendations. The recurring team and its records are invented for teaching. The practical examples include sample instructions and authored illustrations of the results a team would ask AI to produce. They are not captured product runs or evidence of measured business results. Separately labeled research examples describe published studies and include sources.

Read with an activity in mind and, when you reach the workbook, bring it to someone who knows that activity well. Six anchor examples follow the main progression; shorter applications show where the same reasoning travels. The application index at the back groups them by what AI contributes.

Implementation

A first implementation gives people something they can inspect together. A briefing or draft makes it possible to discuss specific choices about the information used and the result expected, rather than relying on a broad description of what AI might do.

The initial example should fit within the permissions and approved environment available to the team. Where those are unclear, resolve them before using organizational information or connecting systems. The aim is to learn enough from a concrete attempt to decide whether the assignment deserves a place in the team's routine.

Choose a Useful Starting Point

We begin with the preparation for Maya's Thursday call. When she shows Lena how it works, she explains that a long summary would not help much. She wants unfinished commitments, changes in the customer's plans, and questions she should resolve with Eli. She also wants to keep the time she spends deciding what the conversation needs. That distinction gives the team somewhere useful to begin.

Look at a real activity

A recent example, with the appropriate records available, will tell you more than a description of someone's usual routine. As the person explains what they did, pay attention to choices that are not obvious in the final result. A customer summary might omit a detail because the customer already resolved it, for instance, or a proposal might contain a question because nobody has confirmed an assumption. A call to a colleague may be how the person settles an ambiguous note. Those choices help you see where AI could contribute and where the person's experience matters.

You might find an opportunity to gather information, prepare a first version, compare alternatives, or help someone think through a question. Keep the conversation broad enough to include better decisions and new possibilities as well as faster preparation.

Maya's team identifies several candidates, including a pre-call brief, a draft follow-up message, and help exploring support for customers who are expanding. They choose the brief first because Maya can explain what she needs and review the result before the customer sees anything. Keeping the other ideas on the list allows them to return to those possibilities once they have learned from an assignment they can examine closely.

Make the contribution understandable

A colleague should be able to understand the proposed contribution without needing a tour of the technology. Maya’s team could describe it as, “Before a customer check-in, prepare a short draft from approved notes so Maya can review commitments and plan the conversation.” The sentence identifies the occasion, the source material, and the person who will use the result. It also makes clear that the result is a draft. “Improve customer experience with AI” can remain the broader ambition, but it leaves too many possible starting points for an initial trial. The narrower description gives the team something it can try and discuss.

ANCHOR EXAMPLE

Find the repeated preparation

The assignment at a glance	
Inputs	Three accounts of the same preparation process
AI contribution	Map repeated steps and propose a candidate use
Human decision	Determine whether the apparent duplication can change
Evidence	Frequency, effort, implementation cost, review effort, and result quality

Lena wants to choose a starting point from several activities. She supplies AI with three short descriptions, gathered from the people doing them:

- **A1, account owner:** Before each check-in, read the last meeting notes, copy open commitments into a preparation document, and ask service colleagues for missing status.
- **A2, service coordinator:** Before the same check-in, read the action list and prepare a status message for the account owner.
- **A3, team leader:** At the weekly review, collect those preparation documents and identify commitments that still lack an owner or next step.

Her instruction is specific:

Map the inputs, repeated steps, and handoffs in A1–A3. Identify one candidate AI contribution and the existing software or human action it would depend on. Do not estimate savings from these descriptions.

An illustrative result identifies two people assembling overlapping status information for the same conversation. A short extract makes the proposed change easier to inspect:

Observation	Candidate change	Dependency
A1 and A2 assemble status for the same check-in	Prepare one source-referenced draft for both colleagues	Customer and action records need stable identifiers
A3 finds commitments without an owner or next step	Flag missing ownership in the same preparation draft	The action list needs a current owner field

Lena checks that map with Maya and Eli. If their descriptions omitted a required review, she adds it before considering a change. AI has helped turn separate accounts of the process into a proposal the team can examine; the people involved determine whether the apparent duplication really can be removed.

An opportunity review can begin this way without pretending it has calculated ROI. To estimate a return, Lena would still need frequency, preparation and review effort, implementation cost, and evidence about the quality of the result.

Choose something you can learn from

For a first assignment, look for an activity that recurs often enough to observe. Someone should understand what a useful result looks like and have time to review it. The necessary information should be available through an approved route.

Also ask whether the team wants the help. A manager may see repetitive preparation where the person doing it sees the moment they begin to understand the customer. Discussing that difference may lead to an assignment in which AI gathers the facts while the account owner keeps

the time spent interpreting them. I would want that distinction settled with the person doing the preparation. Their experience can tell us which effort helps them think and which effort simply makes it harder to get started.

Record the activity, the person who understands it, and why the contribution is worth trying. For Maya, the question is whether an AI-prepared brief helps her arrive ready for a more useful conversation.

Workbook: Choose a useful starting point.

Give AI an Assignment

“Help us take better care of our customers” describes a worthwhile ambition. Maya needs something she can use before Thursday’s call.

Her proposed assignment is a one-page draft drawn from approved customer notes and the open-action list. It should help her find unfinished commitments and changes in the customer’s plans, as well as suggest questions she can consider for the agenda. Maya will review it before the call, so the format needs to make that review straightforward.

Describe the result before building the process

An example of the desired result will help the team settle details that a general request leaves open. In this case, factual points should refer to their source records, and suggested questions should be labeled so they do not read like things the customer has already said. Open commitments need the recorded owner and due date when available, with a clear indication when a date has not been agreed. Conflicting records should be brought to Maya’s attention.

These requirements give whoever configures the AI a target to test. Maya and Eli can review the proposed format together and explain what they would need to check it. Maya may want the customer’s exact question preserved, for example, while Eli may need the action-list reference beside each commitment. It is easier to make those changes while the team is still deciding on the format.

The team’s assignment brief

Part of the assignment	Agreement
Purpose	Help Maya prepare for useful customer check-ins and follow through on previous commitments.
Start	Maya requests the brief before a selected check-in and supplies the approved records.
Inputs	Relevant customer notes and the current open-action list, selected by Maya and used in the approved AI environment.
AI contribution	Prepare a one-page draft with relevant updates, open commitments, and suggested questions. Identify the source record for factual points.
Permitted actions	Prepare the draft. Customer contact, changes to account records, and new commitments remain with people under the existing process.
Human owner	Maya reviews the brief and chooses the agenda. Lena arranges coverage when Maya is away.
Questions	Flag conflicting records or missing information. Ask for unavailable inputs rather than assume customer-specific facts.
Handoff	Maya places the draft in the approved preparation location and marks it reviewed after checking it.
Quality	Important commitments are represented accurately. Facts can be checked. Suggestions are distinguishable from recorded information.
Trial	Try the assignment on five selected check-ins, record what happens, and meet to revise it before expanding use.

Five check-ins is a manageable first trial for this fictional team, not a general threshold for reliability. The first version can be largely manual: Maya selects the records, requests the brief, and saves it for review. This lets the team examine the contribution before deciding whether more of the process should be automated.

Writing the agreement does not configure the software. The prompt guides the result; the software's permissions determine which records and actions are actually available. Whoever sets up the AI needs to make those permissions match a draft-only assignment. Broader actions require their own decisions and setup.

When the assignment continues across steps

The first version ends when Maya receives a draft. A later version might begin from a scheduled check-in, retrieve the permitted notes and action records, notice that the walkthrough status is missing, and wait for Eli's answer before revising the brief. The assignment has become connected, but its purpose has not changed.

Before building that version, the team needs a clearer agreement about movement. The start event must identify the correct account. Each connected system must return a result the next step can interpret. The workflow needs a place to save unfinished work, a route for questions, a cutoff, and a defined final state. Any action outside the preparation workspace remains separately authorized.

Movement question	Agreement for Maya's brief
What starts the work?	A scheduled check-in with a valid account identifier and no existing preparation job.
What may the system read?	Approved notes and open-action records for that account.
What can change the next step?	Retrieved evidence, a failed check, Eli's answer, Maya's decision, or a cutoff.
Where can it pause?	After saving the open question and current evidence for the same preparation job.
What may it write?	One reviewed brief in the internal preparation workspace, after a final authorization check.
What counts as complete?	The saved brief can be found, required sections are present, factual statements have sources, and the run record contains a receipt.

This agreement still does not choose a particular product. It gives an implementation partner something more useful than “make the process

agentic.” The team can first test each step with visible handoffs, then decide where AI judgment is helpful and where ordinary software should keep doing the work.

ANCHOR EXAMPLE

Prepare the customer check-in brief

The assignment at a glance	
Inputs	Approved customer notes and the open-action list
AI contribution	Prepare a source-referenced draft separating facts, unresolved items, and suggested questions
Human decision	Choose what belongs in the conversation and resolve internal questions
Evidence	Corrections, missing items, review effort, and contribution to preparation

Suppose Maya supplies these fictional records:

- **May 6 customer call.** The customer is considering moving staff training from June to July. No new date has been agreed.
- **May 8 action list.** Eli owns a follow-up about a walkthrough. No due date or completion status is recorded.
- **May 9 account note.** The customer asked what support would be available for two possible additional locations. No expansion decision is recorded.

Maya can put those records beneath this instruction:

Prepare a customer check-in brief using only the supplied records. Separate recent updates, follow-up to check, and suggested questions. Identify the source for each factual point. Preserve uncertainty about dates and expansion. Do not treat an absent completion status as proof that an action is unfinished. End with anything I should resolve internally before the call.

The assignment turns mixed notes into a consistent draft while preserving distinctions that matter to the conversation.

A useful draft might read:

Customer check-in brief — draft for review

Recent updates

The customer is considering moving training to July. The date is still open. Source: May 6 customer call.

The customer has asked about support for two possible additional locations. The notes do not record a decision to expand. Source: May 9 account note.

Follow-up to check

Eli owns the walkthrough follow-up. The action list does not show a due date or whether it is complete. Source: May 8 action list.

Suggested questions

Has the preferred training month changed?

What would the customer need to understand before deciding about additional locations?

Before the call

Confirm the walkthrough status with Eli.

The brief keeps the customer's possible expansion separate from an agreed decision and points Maya toward a question she can resolve internally. She can use that preparation while deciding what belongs in the conversation. She may know that training can wait or that a colleague should join the expansion discussion, neither of which would be settled by the draft alone.

Related application: Draft a scope summary from agreed fields

Maya supplies a structured record for two remote training sessions. It identifies the audience, what is included, that on-site support is outside the scope, and that dates and pricing have not been agreed. She asks AI for a short customer-facing summary without adding benefits, prices, dates, or terms.

The draft states the agreed service and preserves the open terms. Maya checks it against the record and the team's approved language before anyone sends it. AI has converted structured information into readable prose; the approved record remains the authority for what is included.

Give the assignment a place

Alongside the intended outcome, an assignment needs operating limits and an agreed response when circumstances change. For Maya's team, that means adding the request and review to the preparation checklist, agreeing how service questions reach Eli, and arranging coverage when Maya is away.

Keep the assignment brief where those people can find it, and revise it when the arrangement changes. Someone using last month's instructions should not have to discover for themselves that the review now belongs to a different person.

Workbook: Give AI an assignment.

Build the First Working Example

A working example lets the team examine whether the assignment produces the result it intended. Maya can prepare a brief from selected records, check it before using it in a customer conversation, and show her colleagues where it needs improvement. They have something concrete to discuss: a missing item, a suggested question, or a format that makes preparation easier.

Begin with an example you can check

The fictional notes in Chapter 2 are enough to begin because Maya can check the draft against a case she understands. They contain a possible schedule change, an open commitment, and a question about additional locations, each of which should appear accurately in the brief.

The review can follow the requirements established for the assignment: identify the source records, separate suggestions from facts, and make open questions visible. When a team moves to real records, those records must be approved for the environment in use, and the reviewer needs to know whether the set is complete. An accurate summary of an incomplete record can still leave someone without important context.

Make the first version easy to inspect

The first version can be manual. Maya selects the inputs, requests the draft, places it in the approved preparation location, and records any changes she makes during review. That provides enough structure to repeat the example and understand how it was produced, without requiring a scheduled trigger or a connection to every account system.

The trial record should include the source references and the version of the assignment instructions used. Otherwise, two people may think they are comparing the same process when one has revised the instructions. If an implementation partner helps configure the system, the assignment

brief gives that person the expected result and the limits on information and actions; the team should be able to see how the setup reflects those limits.

Review the whole attempt

Reviewing the whole attempt includes the effort to find and select records, along with the quality of the draft. A strong output may still need a format change or a simpler way to gather approved inputs before the assignment fits comfortably into the day.

Maya can explain which parts she could use and where she exercised judgment. The first example shows whether the assignment meets its requirements; it cannot establish that every customer conversation will improve.

CONSTRUCTED REVIEW EXERCISE

Correct the brief, then the instruction

Suppose the source note says, “The customer is considering moving staff training from June to July. No new date has been agreed.” An authored imperfect draft turns that into, “Training is scheduled for July.” This is a constructed review exercise, not a captured product run.

Maya traces the statement to the note and changes the draft to, “The customer is considering July; no date is agreed.” She also removes a sentence about last year's seat count. The sentence is accurate, but that question was resolved months ago and has no bearing on Thursday's conversation.

Review finding	Maya's action	Change to the shared instruction
A possible month became a commitment	Restore the uncertainty and retain the source	“Preserve words such as considering, proposed, and not yet agreed.”
An accurate fact is no longer relevant	Remove it from this brief	“Include older context only when it affects an open commitment, current decision, or suggested question.”

The first correction addresses factual fidelity. The second depends on Maya's understanding of the account and the purpose of the call. The team adds both cases to its representative checks, so the review changes the assignment rather than disappearing into one corrected document.

Related application: Answer a question from an internal procedure

A colleague needs to know whether a customer's training can be moved without a new approval. Give AI the question and this small fictional procedure set:

- **P1, date changes:** A coordinator may change the date of an existing remote session after checking trainer availability and receiving written customer confirmation.
- **P2, delivery changes:** A change from remote delivery to on-site delivery requires the service lead's approval.
- **P3, recording:** Record the confirmed date in the action list and retain the customer confirmation reference.

The colleague asks:

The customer wants to move a remote session from June to July and keep it remote. Using P1–P3, explain the next steps and whether service-lead approval is required. Cite the relevant procedure. Say if the answer is not covered.

An illustrative answer says the coordinator can make this date-only change after checking availability and receiving written confirmation, then update the action list. It cites P1 and P3 and explains that P2 applies if the delivery method changes.

The person brings a situational question, and AI selects the applicable instructions and explains how they apply. A shared internal assistant would need access to the current procedure library and a way to return the relevant passages. Microsoft describes this retrieval-and-generation pattern in its current technical documentation. [*RAG and Generative AI*, updated August 4, 2026.](#)

Procedure check · One covered and one uncovered question

The exercise is useful only if it tests restraint as well as retrieval. Keep P1–P3 visible and try these two questions with the same instructions and configuration.

Covered case

The customer wants to move a remote session from June to July and keep it remote. Using P1–P3, explain the next steps and whether service-lead approval is required. Cite the relevant procedure. Say if the answer is not covered.

A passing answer cites P1 and P3, requires trainer availability and written customer confirmation, and explains that P2 applies only if the delivery method changes. It does not invent a date or claim that approval has already been given.

Uncovered case

The customer also asks for a 15% discount because the date changed. Using only P1–P3, state what the procedure allows and what must be referred to someone else. Cite any relevant procedure. Do not infer a pricing rule.

P1–P3 do not establish discount authority. A passing answer may explain the date-change steps that are covered, but it must return the pricing request to an authorized person without inventing a rule. A polished discount explanation is still a failure if the source set does not support it.

Check	Covered case	Uncovered case
Relevant source identified	P1 and P3; P2 only as a delivery-change contrast	No pricing source; P1 and P3 may still support the date change
Required uncertainty preserved	Availability and written confirmation are still pending	Discount authority is unresolved
Prohibited leap avoided	No new date or completed approval	No invented pricing policy or approval
Human handoff	Coordinator completes the documented steps	Authorized pricing owner receives the request with the source gap visible

Record the source version, instructions, system or model, both outputs, corrections, and reviewer. Then change one input: ask for on-site delivery, omit the written confirmation, or supply a conflicting procedure. The point is to learn which distinctions the assignment preserves and where the review needs another representative case. The workbook provides a blank validation record for repeating this test.

Related application: Consolidate records without hiding a conflict

A shared source of context begins with knowing what each record means. Suppose a spreadsheet lists account C17's "training owner" as Robin, while a later note with the same account identifier says, "Rowan will review training; Robin handles access."

An illustrative AI result extracts Rowan as the training reviewer and Robin as the access contact, while flagging that the spreadsheet's "training owner" may describe a different role or may need updating. The account identifier joins the records through ordinary data processing. AI interprets the differently worded roles; a coordinator resolves the conflict and approves any update. The brief and planning comparison should then use the resolved record with its references rather than asking AI to improvise from inconsistent copies.

Decide what is ready to try with the team

The examples may give the group enough confidence to continue with the assignment, or they may reveal a reason to revise it or choose another starting point. If the team continues, its next trial needs named participants, agreed information boundaries, a reviewer, and a date to discuss the results.

Maya can now bring Lena and Eli the source records, draft, and review steps, along with an explanation of what other account owners would need. Making it a shared routine will require time to learn and agreement about what may change around it.

Workbook: Build the first working example.

PART II

Enablement

As more people begin using an assignment, they need to know what they may change, how to carry out their parts, and where questions go. Leaders need a way to hear what the team is learning and act on it, including decisions about time and support.

Enablement begins during the first experiment and continues as AI takes on new responsibilities. The chapters that follow examine the conditions that allow a team to make an assignment part of its work.

Decide What Can Change

Imagine Lena tells the team to try AI, but everyone must continue preparing the old report, use the same meeting agenda, and fit the experiment around a full schedule. Although they have permission to use the tool, the team has very little room to change how it prepares for customer calls.

Maya needs to know whether she can alter the preparation checklist and who can approve a better format. Eli needs to know that helping with the trial is part of his responsibilities. Lena can settle those questions before a wider trial begins, so the team has the authority and time to act on what it learns.

Set a useful scope

Lena brings Maya and Eli together to discuss what they learned from the first example and what a team trial would require. They agree that Maya can change her preparation checklist for the trial. Eli will help clarify open commitments. The existing process for customer communication stays in place. They also agree that the trial will use an approved AI environment and selected records. Any request for additional access goes through the person who manages those permissions. Lena will make room in the schedule for practice and a review of the first results.

The agreement gives the team a defined set of decisions it can make, while identifying those that require another conversation. A scope note could record it:

We will try an AI-prepared customer brief before selected check-ins. Maya may revise the preparation format and will review each brief. Eli will help resolve questions about service commitments. Lena will provide time for the trial and decide whether the team adopts the revised routine. Customer communication and account changes remain under the existing approval process. Additional system access requires separate approval.

A local agreement can be brief while still respecting the policies and obligations that apply to the work.

Some assignments need more than local permission, especially when a trial reaches additional systems, affects employees, or changes customer and financial records. Bring the people responsible for those obligations into the design. Depending on the organization, that may include operations, IT, security, privacy, HR, legal, or a client-services leader; one person may cover several areas in a small company.

They can help choose appropriate records, retention, identity and permissions, and the exceptions that need a person. That work belongs in design, since final approval cannot repair an assignment built around records or authority the team never had.

ANCHOR EXAMPLE

Prepare a new client for kickoff

The assignment at a glance	
Inputs	Signed agreement, intake form, kickoff email, and onboarding checklist
AI contribution	Reconcile records, identify missing information, and draft a clarification and agenda
Human decision	Verify the intake, approve communication, and decide when the record may advance
Evidence	Extraction corrections, unresolved fields, review effort, and duplicate work removed

A signed agreement reaches an onboarding coordinator. The agreement says the customer is purchasing two remote training sessions; the intake form names an account contact but leaves the technical contact blank. A kickoff email says, “Please include our new location,” without identifying which location. The checklist requires the agreed service, account and technical contacts, location, training audience, and desired outcomes before kickoff.

The proposed AI assignment compares the agreement, intake form, and email with the onboarding checklist. Its instruction is:

Extract the agreed service and named contacts. Match them to the required intake fields. Identify missing or conflicting information, with its source. Draft a clarification message and a proposed kickoff agenda. Do not create users, expand scope, or send anything.

The expected result includes the two agreed sessions, a missing technical contact, and a location question linked to the email. Its clarification draft asks the customer to identify the location and technical contact; the agenda includes confirming the training audience and desired outcomes.

AI is interpreting several forms of language and turning them into a consistent intake record. The customer record's creation, required-field check, and kickoff booking can be ordinary workflow steps. If account provisioning is included later, it should use approved roles and verified identities rather than let the model decide who gets access.

The coordinator checks the extracted fields and raises a concern: some remote engagements do not need a separate technical contact. Making that field universally mandatory would create another unnecessary task. Lena revises the rule so the field must be completed or marked “not needed” with a reason. She assigns checklist maintenance to the onboarding coordinator and confirms that the reviewed intake replaces the separate kickoff worksheet during the trial.

When the customer's answers arrive, the coordinator approves the record for the next stage. The AI output has changed the intake route, but the enablement decision is visible in the retired worksheet, the revised rule, and the named owner.

Make participation possible

Enthusiasm doesn't create an opening in someone's calendar. Learning a new routine takes time, whether that comes from part of an existing team meeting or postponing a lower-priority task. Lena's scheduling decisions show that practice and feedback belong among the team's responsibilities.

Ethan Mollick distinguishes individual use from organizational practice and gives leadership, experimentation, and employee participation explicit roles. [“Making AI Work: Leadership, Lab, and Crowd,” May 22,](#)

2025. For our team, participation means Maya can say the brief is too long, Eli can explain an ambiguous action, and Lena can change the arrangement in response.

Be willing to revisit a responsibility

Examining an activity may reveal that it has no clear owner or that two people gather the same information. A checklist may resolve a local handoff. A lasting responsibility, such as Eli maintaining information the brief depends on or Maya helping colleagues use the process, needs recognized time and backup.

I would let the need determine the size of the change. An org-chart redesign may become appropriate, but a team can begin by adjusting schedules, ownership, and the point at which colleagues work together.

Tell people what success will mean

Lena wants account owners to be prepared and able to give customers more attention. She can explain that the trial will help the team learn whether this assignment supports that aim, and how she intends to act on the feedback. A brief that takes too much time to review may need revision. One that helps account owners prepare may be worth adopting, with additional support if needed. Knowing how their observations will inform those decisions gives participants a reason to report honestly about a practice they may continue using.

Before the wider trial starts, each participant should be able to answer a simple question about authority: “If I learn that something needs to change, who can decide?” Write the answer in the scope note. It will save the team from having to rediscover it during the first busy week.

Workbook: Decide what can change.

Help People Work Well With AI

Maya understands the brief because she helped design it. The next account owner may need help selecting notes, checking a factual point, or deciding whether a suggested question belongs in the customer agenda. That colleague may also see an application Maya has not considered, giving the team another opportunity to improve the assignment as people learn it.

Ethan Mollick's *Co-Intelligence* considers AI in both work and learning, including its role as a collaborator and coach. A team can bring that perspective into practice sessions, with people working alongside AI and helping one another interpret and improve what happens.

Co-Intelligence.

Practice with an example people recognize

A practice session can use a fictional or appropriately approved example that resembles the team's activity, with the source records and draft beside each other. Participants check a commitment against its source, identify a suggestion that needs human judgment, and explain what they would change before using the brief in a conversation. Those actions give "become comfortable with AI" a practical meaning. A participant should be able to prepare the inputs, request a draft in the approved environment, review it, and explain what happens next. After watching a demonstration, they can try the process themselves with room to ask about the work as well as the software.

A question about including a customer's tentative plan or an older note may lead the team back to the constructed review in Chapter 3. Practice is also a chance to refine the assignment using the experience of the people who will carry it out.

Use uncertainty as a practice exercise

The fictional records in Chapter 2 let a learner practice distinguishing a possibility from a commitment. After checking the brief, the learner can rehearse the conversation with AI.

Related application: Rehearse the planning conversation

Maya supplies the three records and asks AI to play the customer. Her instruction describes the exercise:

You are a customer considering two additional locations. No expansion decision or training date is agreed. Ask one question at a time. Keep your answers consistent with these facts. After three exchanges, give feedback against these criteria: did I clarify your decision, distinguish options from commitments, and agree a next step? Refer to words I used. Do not invent company policy.

A short illustrative exchange might go like this:

AI as customer: Could you get both new locations trained in July?

Maya: Before we look at dates, have you decided to open both locations, or are you still comparing options?

AI as customer: We're still comparing options. I need to understand the training implications first.

Maya: Then let's identify who would need training and who could support it. Eli and I can use that to prepare options before we discuss a date.

AI can identify Maya's clarifying question and her decision to establish who needs training before discussing a schedule. Suppose it also praises her for agreeing on a next step. The facilitator points out that Maya proposed preparing options but did not agree on who would send them or when. The feedback was partly useful, and interpreting it became part of the exercise.

The learner can repeat with a customer who has different questions without requiring a colleague to perform every rehearsal. The exercise supports practice; it does not certify competence.

RESEARCH EXAMPLE

Assistance during actual customer-support work

In *Generative AI at Work*, Erik Brynjolfsson, Danielle Li, and Lindsey Raymond studied an assistant that supplied response suggestions during customer chats. The human representatives could edit or ignore those suggestions. Using data from 5,172 support agents, the published study found a 15% average increase in issues resolved per hour, with larger benefits among less experienced workers. These are results from one firm's deployment, not a general savings promise. They show a concrete way AI support can become part of the activity itself. **The Quarterly Journal of Economics**, 2025.

Let leaders participate

Trying the process herself gives Lena a basis for deciding what support the team needs. She may find that selecting records takes longer than expected or that the review instructions need an example. Asking questions and improving the instruction with colleagues also makes practical knowledge welcome, regardless of who brings it.

Give everyone a way to contribute

People can help shape an AI-supported activity without becoming software developers. An account owner may describe the result she needs and the customer distinctions the draft must preserve. A coordinator may identify the exception that should stop an automated step. Someone who has never configured an agent may still be the person who can tell whether its work belongs in the routine.

The team should make those contributions easy to record and act on. Some colleagues may choose to build a small example; others may test one, supply representative cases, question a rule, or improve the handoff.

Production software still needs named responsibility for access, testing, maintenance, and support. Calling everyone a builder can obscure that difference and leave the maintenance work with whoever was most enthusiastic at the beginning.

Support different starting points

Some people will explore immediately; others will prefer a demonstration and a familiar example. A short guide and a colleague to ask can support both, but avoid making the most enthusiastic person the permanent help desk by accident. If someone supports colleagues, give that responsibility an owner, time, and backup.

A report of a confusing result should identify the inputs, the expected result, and what appeared. That gives the team something specific to investigate without asking the participant to defend the decision to try AI.

Turn feedback into a visible improvement

Suppose a participant finds that the brief includes too much history. They want open commitments near the top and older context only when it affects the next call. Maya and Eli can review that suggestion and try a revised format on another case, then update the shared instructions if they adopt the change.

Showing colleagues how their feedback affected the brief gives them a reason to contribute again. As each participant works through an example with less help, the team can see which steps are becoming familiar and which deserve another demonstration.

Workbook: [Help people work well with AI.](#)

Connect the Responsibilities

Maya opens the brief and sees that she needs to check the walkthrough status with Eli. The question has surfaced in time to help her prepare, provided it reaches him and his answer becomes available for the call. If Maya learns that the walkthrough already happened, the team needs a place to record that answer so the same unresolved item does not reappear next time. The handoff includes both getting a response and keeping the shared record current.

Trace one item all the way through

Following a single commitment through the process can reveal where a handoff needs attention. In our example, Eli coordinates the customer’s walkthrough, and Maya needs its current status before the next conversation. She needs to know where the request is recorded and how to find his update; he needs a reliable way to receive the request and record what happened.

The action list may need a clear completion field, or Eli may have been recording answers somewhere Maya doesn’t normally look. Agreeing on one place for the current status gives both people a record they can rely on and makes future AI assignments easier to design. Here is the first version of the handoff for this fictional team:

Stage	Who acts	What the next person receives
Prepare inputs	Maya	Selected approved notes and the current action list.
Draft	AI, when Maya requests it	A source-referenced brief with questions marked.
Resolve service questions	Maya asks Eli; Eli responds	Confirmed status and an updated action entry where needed.
Review	Maya	A checked brief and a chosen agenda.

Stage	Who acts	What the next person receives
Meet with the customer	Maya	New notes and agreed next steps.
Record commitments	Maya and the relevant service owner	Updated actions with owners and dates when agreed.

Naming Maya and Eli at the relevant stages makes each person’s responsibility clearer than an instruction that “the team reviews it.”

Maintain the context the assignment depends on

The brief depends on more than a collection of documents. It needs stable account identifiers, a shared meaning for an open commitment, current service rules, and a known record for the latest status. Those definitions and responsibilities form part of the context in which AI works.

A team does not need to move every record into one large system before it can begin. It does need to know which record is authoritative, who keeps it current, how the AI reaches it, and what happens when two sources disagree. A connector can retrieve information, but it cannot settle unclear ownership or repair a business rule that nobody maintains.

Useful context does not require recording everything, so keep the records that support the activity, learning, and accountability within the agreed purpose. Unnecessary capture gives the team more sensitive material to protect and more irrelevant material for AI to retrieve. Maya's team needs the customer commitments and the result of the review; it does not need an indiscriminate recording of every conversation around the account.

Separate a draft from a decision

Suppose the brief suggests asking the customer about support for additional locations. Maya may choose to discuss it, and the customer may make a request that requires someone else’s decision. As that information moves, people need to distinguish an idea under consideration from a commitment to provide a service or meet a date.

Simple labels such as “draft,” “reviewed,” and “agreed action” can help, provided the team understands their meanings. “Reviewed” might mean

Maya checked the brief against the supplied records; it should not imply that every suggestion has been approved for action. With that meaning established, colleagues can use the document without guessing whether it contains a proposal or a commitment.

Give questions somewhere to go

The manual brief gives Maya a question to resolve with Eli or Lena. An agent-based arrangement can make that handoff part of the assignment, provided the system has the permitted tools and a way to wait for an answer.

Related application: Handle an incoming support request

A customer writes, “Our new coordinator can't get into the training portal, and could we move next week's walkthrough?” The approved support article says users should use the portal's password-reset page; staff must not request passwords. Scheduling requires a confirmed customer choice from available slots.

The AI assignment identifies two requests in the message: account access and rescheduling. It can retrieve the approved access article and draft a response with the reset instructions, while separating the scheduling request for a calendar lookup. A sample response for review is:

Please use the password-reset option on the training portal. We won't need your password. I can also help find another walkthrough time; I'll check availability before offering alternatives.

The contribution is recognizing both needs in an unstructured message, finding the relevant guidance, and preparing the next step. A routing rule sends the access request to the support queue, and a calendar service returns real availability. A coordinator reviews the response during the trial. If the approved material does not cover the problem, the request goes to a person with the original message attached.

Connect more steps when the team is ready

For an agent-based process, the connections need to be explicit enough to implement and test, including what happens when a person or system does not respond. The next useful step may depend on what a tool returns, which makes the assignment different from a fixed chain of screens or notifications.

Three kinds of movement

Three related ideas are easy to collapse into the word *workflow*. Keeping them separate helps a leader ask better questions about the design.

Idea	What moves	Question to ask
Loop within a run	AI chooses a permitted next step, uses a tool, examines the result, and continues, pauses, completes, or stops.	What result changes the next step, and what ends this attempt?
Workflow graph	Configured routes show where work branches, joins, waits, or returns for another attempt.	Which routes are fixed in software, and which choices rely on AI judgment?
Improvement between runs	People use observed failures to revise instructions, examples, tools, or checks, then evaluate the new version.	What changed after the last attempt, who accepted it, and did the revision help?


The loop happens during one attempt. It may retrieve a record, discover a gap, and choose a targeted lookup before it can draft. The graph is the larger route that permits that movement and places limits around it. Improvement between runs happens later, when the team changes the assignment itself.

A graph does not automatically make a process agentic. A routing rule can send every missing status to the same queue without any model deciding what to do. An agent can also operate inside one node of a mostly deterministic graph, choosing among a small set of permitted lookups. Multiple agents are optional. Dividing a task among several

models may be useful, but the defining question is still how the next step is chosen and controlled.

OpenAI's Agents SDK documents a runner loop in which tool calls are executed, results are appended, and another model call follows until final output or a turn limit. Anthropic uses a similar distinction between predefined workflows and systems in which a model directs its tool use. These are available construction patterns, not evidence that a connected business assignment will finish correctly. [OpenAI, "Running agents," accessed September 6, 2026.](#) [Anthropic, "Building effective agents," December 19, 2024.](#)

Figure · The execution loop within one run



A result may send the same run to another permitted step, a person, completion, or a controlled stop.

The figure begins from an assignment and available evidence. A tool result can send the loop toward another permitted step, a person, completion, or a controlled stop. The loop should not continue merely because another action is possible.

A capability map for leaders

The map below describes mechanisms that current agent frameworks and retrieval systems can support. Availability depends on the model, runtime, connected tools, permissions, and application design. A chatbot with no configured tools cannot retrieve an account record simply because another product can.

Capability	Business contribution in this case	Setup the team would need
Tool use and artifact creation	Read an action record, place a question in an internal queue, or save an approved brief	A defined tool, structured inputs and results, scoped credentials, error handling, and authorization in the connected system
Targeted retrieval	Search for the current walkthrough status after the first records prove insufficient	Searchable approved sources, identifiers, access filtering, ranking, source references, and a response for no useful result
Branching	Choose whether to draft, ask Eli, revise, complete, or stop	Explicit states, allowed routes, criteria for each route, and a default route when the result is ambiguous
Parallel checks	Retrieve notes and actions at the same time, or run separate factual and required-field checks	Independent tasks, join logic, timeout behavior, and a rule for partial results
Iterative revision	Return an unsupported statement for one bounded correction	Check criteria, usable feedback, a revision limit, and a failure state when the check still does not pass
Saved progress	Wait for Eli without rebuilding the whole preparation from memory	Durable state, a stable run identifier, retention rules, version compatibility, and a tested resume path
Bounded action	Write one approved brief and verify that it exists	Narrow permission, approval where required, duplicate prevention, post-action verification, and a correction or reversal route

OpenAI's tool documentation shows that an agent runtime can expose data retrieval, code, files, and application functions as tools. It also states that hiding or showing a tool does not replace authorization inside the tool or connected service. Microsoft documents both classic retrieval and an agentic retrieval option that plans focused subqueries and runs them in parallel; that option requires configured knowledge sources and a

knowledge base. [OpenAI, “Tools,” accessed September 6, 2026.](#) [Microsoft, *RAG and Generative AI*, updated August 4, 2026.](#)

Parallel work deserves the same care as sequential work. Two reads can finish in either order, one may time out, and a join must know whether a partial result is acceptable. Running several model checks can expose different concerns, but agreement among models is not independent verification. A source lookup, calculation, system receipt, or qualified human judgment may still be the evidence the assignment needs.

What the system keeps

The word *memory* can refer to several different things. A team should name the one it means before deciding what to store.

Term	Working meaning in this book	Example
Context	Information supplied or retrieved for the current model step	The check-in date, customer notes, action records, and current instruction
Checkpoint	A saved snapshot of one run's state so that the same attempt can resume	Evidence gathered, the open walkthrough question, the next allowed routes, and the run identifier
Stored memory	Selected information retained for later interactions or other runs	A confirmed customer preference or a reusable lesson, stored under an agreed purpose and retention rule
Reusable instructions	Versioned directions, examples, schemas, and checks applied to future runs	“Preserve words such as considering and not yet agreed,” plus the required brief sections
Model training	A separate development process that changes a model using a training dataset and evaluation	A fine-tuning job, if the organization later has enough justified examples and a reason to create a different model

Saving a conversation is not the same as training a model. A checkpoint does not teach every future run what happened, and a reusable instruction can change behavior without changing the model's

parameters. LangGraph's current documentation distinguishes thread-scoped checkpoints from application-defined information stored across threads. OpenAI's fine-tuning API describes training as a separate job that creates a model from an uploaded dataset. [LangGraph, "Persistence," accessed September 6, 2026.](#) [OpenAI, "Fine-tuning," accessed September 6, 2026.](#)

For Maya's brief, the team needs a checkpoint while it waits for Eli. It does not need to convert every customer conversation into long-term memory. Any fact kept for another run should have an identified source, purpose, owner, and retention period.

Let responsibility grow with evidence

The first version of an assignment can remain visible and reversible. As the team learns, it may allow AI to take on more of the surrounding steps. The progression below is one way to make that decision explicit.

Responsibility level	AI contribution	Human responsibility	What the team needs to see before expanding it
Draft	Prepare a result from selected records without changing another system	Review the result and decide what to use	The draft can be checked against representative cases and the reviewer can identify errors efficiently
Recommend	Identify a condition and propose a next step	Choose whether the recommendation fits the situation	Recommendations are relevant, important exceptions reach the right person, and rejection reasons can improve the assignment
Prepare an action	Complete the permitted preparation and place an action in a review queue	Approve, revise, or reject the action	Permissions, duplicate handling, cutoff times, and recovery from a failed step have been tested
Act within a boundary	Take a defined action and record what happened	Set the boundary, review exceptions, and examine the continuing results	The action has dependable checks, a clear reversal or correction path, and evidence that the review arrangement remains sufficient

The last row is not the required destination, since a customer commitment or employment decision may always remain with a person while a routine record update could eventually be allowed within a narrow boundary. The team should choose the level that fits the consequence of the action and the evidence it has gathered.

Human judgment may also appear throughout the activity rather than only at its beginning and end. Maya selects the relevant situation, Eli resolves an ambiguous status, the coordinator maintains the checklist,

and Lena changes the surrounding responsibility. Removing any one of those contributions would change the assignment itself.


ANCHOR EXAMPLE

Follow one briefing from trigger to receipt

The assignment at a glance	
Inputs	Scheduled check-in, authorized notes and actions, lookup results, and Eli's dated answer
AI contribution	Choose targeted retrieval, identify a gap, draft, and revise against a check
Human decision	Supply missing status, judge the agenda, approve the brief, and retain authority for customer commitments
Evidence	Source references, route choices, pauses, corrections, completion checks, write receipt, and review effort

The assignment is still narrow: prepare Maya's check-in brief from authorized account records, resolve one missing walkthrough status when possible, and write the accepted brief to the internal preparation workspace. It does not contact the customer or change a commitment. The trace below is an authored teaching example, not a captured run from installed software.

Figure · The briefing workflow and its decision points



Maya's briefing route combines fixed software, AI judgment, human decisions, saved state, and a bounded revision loop.

The graph uses three kinds of nodes. Ordinary software validates identifiers, starts parallel reads, enforces limits, and verifies writes. AI judgment selects relevant passages, identifies the unresolved status, chooses a targeted lookup, drafts, and revises. People supply facts, decide whether the brief is useful, and authorize consequential action.

The successful run

Step and type	What happens	Result that changes the next step
1 · Ordinary software	A calendar event identifies account C17. A duplicate-prevention check finds no existing preparation job and assigns run C17-04.	A valid account and new run allow the two reads to begin.
2 · Ordinary software, parallel	The workflow retrieves permitted account notes and open actions at the same time. Each result retains record IDs and dates.	Both reads succeed, so the workflow joins their evidence.
3 · AI judgment	AI selects the records relevant to the coming conversation and notices that the walkthrough has an owner but no current status.	The missing status makes another lookup useful before drafting.
4 · AI judgment plus retrieval	AI requests a targeted search for records tied to the walkthrough identifier. The result is an older scheduling note, not evidence that the walkthrough occurred.	The gap remains, so the assignment routes one precise question to Eli.
5 · Ordinary software	The workflow checks that no question for this run already exists, places one item in Eli's internal queue, saves the evidence and open question, and pauses.	The saved checkpoint waits for Eli's answer or the review cutoff.
6 · Human decision	Eli records, "Walkthrough completed May 12; follow-up guide due May 16," with the service record reference.	The dated answer resumes run C17-04 and becomes part of its evidence.
7 · AI judgment	AI drafts the brief with recent updates, follow-up, suggested questions, and source references.	The draft goes to factual-support and required-field checks.

Step and type	What happens	Result that changes the next step
8 · Check and bounded revision	The draft says the customer “will add two locations,” although the May 9 note says only that the customer asked about support for two possible locations. The check returns that statement with its source. AI revises it to preserve the possibility.	The second check passes. The revision count is one of two allowed attempts.
9 · Human decision	Maya checks the sources, removes one question that does not belong in the conversation, and approves the brief for the preparation workspace.	Her approval permits the internal write; it does not approve customer contact or a new commitment.
10 · Ordinary software	The write step rechecks the destination and Maya's approval, uses run C17-04 as its duplicate-prevention key, saves the brief once, reads it back, and records the document ID and time.	A matching receipt satisfies the final completion check.

The tool result changes the next step more than once. The old lookup result sends the run to Eli. Eli's answer sends it back to drafting. The unsupported statement sends it through a bounded revision loop. Maya's approval permits one internal write. No second agent is necessary to make those branches possible.

The interrupted run

Run C22-02 reaches the same walkthrough-status gap, but Eli does not answer before the review cutoff. The workflow keeps the evidence it has, marks the question open, and returns an incomplete brief to Maya. It does not convert the absence of an answer into “not completed,” and it does not keep calling tools after the cutoff.

Completion state	What Maya receives	What the system must not do
Complete	A reviewed brief with required sections, supported factual statements, an internal document receipt, and a run record	Claim that the customer meeting or follow-up action has happened
Stopped incomplete	A draft marked incomplete, the unresolved question, sources gathered, and the reason for stopping	Invent the status, hide the missing evidence, or write to the accepted-brief location
Failed	A visible error, the last confirmed checkpoint, and the route for retry or human handling	Repeat an uncertain external action without checking what already happened

The stopped run can resume if Eli answers while the preparation is still useful. If the customer meeting has already occurred or the instructions changed materially, the owner decides whether to start a new run rather than quietly resume old work under new conditions.

Recover after an uncertain action

Suppose the final write request times out. The system cannot tell from the missing response whether the document was never saved or saved successfully. Retrying immediately could create a duplicate. The recovery route first looks for a document carrying the run key C17-04 in the intended location. If it finds one with the expected content hash, it records the missing receipt and completes. If it finds none, it retries once. If it finds a conflicting document, it stops for a person.

A checkpoint preserves the workflow's state, but it does not make an external action safe to repeat. The write operation needs its own duplicate-prevention key, result check, and correction route. LangGraph's interrupt guidance makes the same engineering point: operations that may run again around a pause should be idempotent or separated into steps that can be verified before repetition. [LangGraph, “Interrupts,” accessed September 6, 2026.](#)

OpenAI's human-in-the-loop documentation describes pausing a run before a sensitive tool call, serializing its state, and resuming after a

decision. LangGraph documents conditional interrupts and thread identifiers for loading saved state. Those mechanisms require deliberate storage and versioning; they do not appear in every chat interface or connector automatically. [OpenAI, “Human-in-the-loop,” accessed September 6, 2026.](#) [LangGraph, “Persistence,” accessed September 6, 2026.](#)

This arrangement gives AI a recurring contribution across several steps. People no longer have to remember to start each draft, repeat the same record search, or carry every question between documents. The remaining human work is clearer too: Eli resolves the service fact, Maya judges the conversation, and Lena decides whether the assignment deserves a wider place in the routine.

Workbook: [Connect the responsibilities.](#)

Build a Practice That Improves

The brief reads well, and the team likes its format. To decide whether to keep using it, they need to examine how it helps Maya prepare and conduct a customer conversation. Their original purpose gives them a basis for that review: make relevant information easier to examine, keep track of previous commitments, and leave more attention available for the customer.

Understand the current approach

A comparison begins with the way people actually prepare, including finding notes, clarifying questions with colleagues, and choosing an agenda. Where time matters to the decision, representative examples can establish how long that preparation takes. Quality also needs a working definition, such as accurate commitments, relevant questions, and source references an account owner can follow. If the trial has already begun without a baseline, record that limitation and gather a current comparison where practical. Do not invent what the previous routine required.

The AI-supported version needs to be examined using the same definitions, including input selection, review, and corrections. Initial setup and training effort belong in the record too. Enough context should accompany each example to distinguish a routine check-in from a complicated account transition, which may require very different preparation.

Test the assignment, not a general promise

A randomized study with 758 BCG consultants examined GPT-4 on realistic consulting tasks. On tasks within the model's tested capabilities, participants with AI completed work 25.1% more quickly on average and produced higher-quality solutions. A separate task outside that range showed worse accuracy with AI. The findings support evaluating the particular assignment and its review requirements, rather than treating one successful task as proof about every task. [Dell'Acqua and colleagues, *Navigating the Jagged Technological Frontier*](#).

Define completion and failure before the trial

A connected assignment can produce a fluent brief while still failing to complete the work. The team therefore defines completion in terms of evidence and state, not whether the last model response sounds finished.

Check	How it is performed	Pass condition	Failure response
Correct account and run	Ordinary software compares the event, account identifier, and duplicate-prevention key	One active preparation run belongs to the scheduled check-in	Stop and route the identifier problem to the workflow owner
Required evidence	Software checks that the expected reads returned; AI identifies whether the content answers the assignment	Required sources are present, or each missing item is visible	Continue to a permitted lookup, ask a person, or stop incomplete
Factual support	A source check compares factual statements with cited records; Maya reviews consequential distinctions	Statements are supported and possibilities remain possibilities	Return repairable language for revision; stop after the revision limit

Check	How it is performed	Pass condition	Failure response
Required sections	A schema or field check looks for updates, follow-up, suggested questions, and sources	Every required section is present, even when one says no supported item was found	Return the missing field once, then fail if it remains absent
Human review	Maya checks usefulness, relevance, and the agenda	Maya approves the brief or records why she rejected it	Revise only if the reason fits the assignment; otherwise stop
Final write	Software checks approval and destination, writes once, reads back, and records a receipt	The expected document exists once and matches the accepted result	Inspect before retrying; stop on a conflict or exhausted retry

Some checks are deterministic. A schema can confirm that a source field exists, a calculation can verify a total, and a system can return a document identifier. Other checks require judgment. Maya decides whether a question belongs in the conversation, and Eli decides whether a service note establishes the walkthrough status. A model evaluator can help find unsupported statements, but agreement between two models does not prove that the underlying record is complete or that the brief is useful.

Notice more than speed

Better preparation can be valuable even when it takes about the same amount of time. Maya might arrive with more relevant questions, Eli might receive earlier notice of a service commitment, or a colleague covering a call might find the customer history easier to understand. These are possibilities to examine through specific examples from the people using the brief.

“It was good” is encouraging feedback. “It reminded me to confirm the walkthrough before the call” gives the team something it can inspect: whether the question was present and whether someone followed it up. That observation may explain how the brief supported a conversation, though claims about improved retention or new revenue would require their own evidence.

Related application: Produce a report that explains the records

Suppose Lena receives these fictional weekly service records:

Record	Planned sessions	Completed sessions	Note
W1	2	2	Both sessions delivered
W2	3	2	Customer moved one session to next week
W3	1	0	Trainer unavailable; replacement not confirmed

A spreadsheet or query calculates six planned sessions and four completed. AI explains that W2 records a customer-requested move and W3 still needs a replacement trainer, linking each statement to its record. The team leader checks the calculation and references before releasing the report. The arithmetic and data refresh should remain deterministic; AI contributes the language-based interpretation. A live version also needs consistent date ranges, fresh sources, and reconciliation between the query and narrative.

Keep a modest trial record

The team can capture the first results without creating a large reporting burden.

What to record	Why it helps
Type of check-in and records used	Makes the example understandable and supports a fair comparison.
Preparation and review effort	Shows what the full routine requires.
Important corrections or missing items	Helps revise the assignment and review instructions.
A useful contribution	Identifies what the account owner could do with the brief.

What to record	Why it helps
Next adjustment	Gives the feedback a practical destination.

Avoid collecting sensitive customer details in a trial log that does not need them. Use an approved reference when that is enough to find the underlying record.

The team’s observations should guide whether it continues with the brief as written, shortens it, improves the input records, or tries another type of check-in.

A completed row for one invented trial case might read:

Check-in	Review effort	Correction	Useful contribution	Next adjustment
Expansion discussion	8 minutes	Restored “considering July”; removed resolved seat-count history	Prompted Maya to confirm the walkthrough and ask what the customer needed before deciding	Add both review cases from Chapter 3 to the representative checks

Count accepted results and the work required

The cost of one model call is not the cost of a result the team can use. A run may include retrieval, several model calls, a correction, human review, and a failed action that has to be reconciled. Runs that stop incomplete still consume resources and belong in the total.

The following numbers are invented to demonstrate the calculation. They are not prices from a provider, measured labor costs, or a forecast for Maya's team. Assume the organization uses an internal review-cost rate of \$36 per hour for this exercise.

Run	Outcome	AI service and tools	Review and correction	Full run cost
R1	Accepted on first draft	\$0.40	5 minutes = \$3.00	\$3.40
R2	Accepted after one revision	\$0.65	8 minutes = \$4.80	\$5.45
R3	Stopped incomplete at cutoff	\$0.30	4 minutes = \$2.40	\$2.70

Run	Outcome	AI service and tools	Review and correction	Full run cost
R4	Accepted on first draft	\$0.35	4 minutes = \$2.40	\$2.75
R5	Accepted after uncertain-write recovery	\$0.70	9 minutes = \$5.40	\$6.10

The five runs cost \$20.40 in total and produced four accepted briefs, for an illustrative accepted-result cost of \$5.10. Reporting only the four successful runs would hide the cost of R3. Reporting only model usage would hide most of the review and recovery effort.

The same calculation can use minutes instead of dollars when the organization does not have an agreed internal rate. Keep the components separate enough to see whether a change reduces service expense by moving work onto the reviewer. Also record the number of runs that stopped, because an apparently low accepted-result cost can coexist with an unacceptable failure rate if the sample is small or selective.

Change the assignment between runs

The correction in run C17-04 happens inside one run. The check identifies an unsupported commitment, AI revises the sentence, and the second check passes. That does not by itself improve the next run.

Between runs, Maya and the workflow owner examine the failure and change instruction version 3. Version 4 adds a required classification before drafting: each statement about customer plans must be marked internally as a recorded commitment, a possibility, a proposal, or unresolved. The brief can omit those labels from customer-facing prose, but the factual-support check sees them alongside the source.

They test both versions on four invented cases:

Representative case	Version 3	Version 4
Customer is considering July; no date agreed	Fails by stating July as scheduled	Preserves the possibility
Customer asked about two possible locations	Fails by stating that two locations will open	Preserves the possibility
Signed order confirms two remote sessions	States the two-session commitment correctly	States the commitment correctly
Action has an owner but no current status	Keeps the status unresolved	Keeps the status unresolved

Version 4 passes these four checks and becomes the candidate for the next limited trial. Four cases do not establish reliability across every account. They show that one specific change addressed the observed failure without breaking three nearby distinctions.

Activity	Happens when	Record to keep
Revision within a run	A repairable output fails a check and the run still has revision budget	Failed statement, feedback, revised statement, check result, and attempt count
Improvement between runs	The owner changes the shared assignment after reviewing one or more runs	Failure pattern, instruction or configuration change, representative cases, comparison result, approver, and new version

Anthropic's current evaluation guidance describes agent evaluation as testing multi-turn behavior against explicit success criteria, including what changes in the environment. That supports examining the run rather than grading only its final prose. The guide is an engineering perspective, not evidence that an automated grader can replace source checks or human judgment. [Anthropic, “Demystifying evals for AI agents,” January 9, 2026.](#)

Related application: Reconcile billing language with delivered work
A billing coordinator has three records for a fictional engagement:

- **S1, approved scope:** Two remote training sessions at \$500 each; additional sessions require approval.
- **D1, delivery log:** Two remote sessions completed.
- **I1, draft invoice:** Three “enablement workshops” at \$500 each.

AI links “enablement workshops” to the remote sessions as a possible terminology match, then flags that the invoice bills three units while the scope and delivery log support two. A calculation shows a \$500 difference. The coordinator verifies the match and finds no approval or delivery record for a third unit, so the invoice is corrected. The team adds the accepted terminology match to its service map and the missing-third-record case to future checks. Exact identifiers and quantity checks still do as much of the reconciliation as possible; AI contributes the remaining language match and explanation.

Ask what people would miss

“If AI disappeared from your organization tomorrow, what would your team miss most?” can open a conversation about the support people value. A colleague may appreciate being able to cover a customer conversation with better context, while another may value having more ideas to consider.

Specific examples can connect those answers to the trial record and help explain AI’s contribution in language the team recognizes. The discussion should also leave people able to describe the activity itself and how they would continue when needed, so appreciation for the support goes alongside an understanding of the work.

Make a decision the evidence supports

At the agreed review, participants can compare the examples with the original purpose and identify what they still need to learn. A decision might be, “We’ll continue for these check-ins, move commitments to the top, and review the next set before offering it to the rest of the team.” That decision names both the next step and its limits. Explaining why the assignment is worth continuing, along with any revisions it needs, gives the team a stronger basis for investing in AI and describing its contribution.

Keep the shared practice usable

Keep the current instructions, examples, owner, status, and last review date in one dependable place. A new colleague should be able to distinguish an experiment from an established practice, find the right version, and see who can approve a change.

Make a run traceable

When an assignment repeats, retain enough information to understand how a particular result was produced. For one of Maya's briefs, that might include the check-in identifier, the source-record references, instruction version, model or configuration, resulting brief, Maya's correction, and the action taken after review. The record can point to approved material rather than copy sensitive content into a separate log.

Suppose Maya removes an old seat-count note from brief C17 because it has no bearing on the current conversation. The run record connects that correction to instruction version 3 and configuration M2. When the team updates the instruction and tests C17 again, it can see whether the older fact stays out without losing a relevant historical commitment, giving the feedback a specific destination.

Alex Lieberman's 2026 practitioner checklist for AI-native companies connects evaluation, graduated autonomy, and traceability in a similar way. It is an operating perspective rather than research evidence, but the combination is useful: broader action can follow specific checks, and the team can understand why a version changed. ["30 Features of an AI-Native Company," August 31, 2026.](#)

Make support part of the arrangement

Supporting an assignment may grow into a regular responsibility. Its owner needs time and backup to maintain examples, answer questions, and help new participants practice. A short change record can help another team distinguish local choices from the parts it can reuse.

Operating cost belongs in the trial record alongside review effort. Cost per accepted result is a more useful comparison than the price of one model call. Include the AI service, connected tools, retries, review time, and corrections required before the result can be used. A cheaper

configuration is worthwhile only if it passes the same checks without shifting additional work to the reviewer.

Decisions about replacing existing software should come later and include maintenance, access controls, support, and record portability. They are separate from showing that AI contributes to the activity.

Revisit the assignment when the activity changes

The owner repeats representative checks when the records, responsibilities, actions, system, model, or configuration change in a way that could affect the result. A stable assignment may need only a scheduled review; a material change deserves closer examination.

A new model is a reason to compare, not a requirement to rebuild a working routine. The team can keep the current arrangement when the alternative does not improve the result, cost, or experience enough to justify the disruption. Review should follow evidence and changes in the activity rather than a calendar that assumes every workflow must be reinvented every few months. Retire the assignment when it no longer serves its purpose.

Workbook: Build a practice that improves.

Transformation

An AI assignment can become a dependable part of the routine and be worth keeping there. It can also give a team a reason to reconsider the service that routine supports. The preparation brief gives Maya and her colleagues an opportunity to have a different conversation with customers, if Lena is willing to make room for it.

Suppose the team has found that the reviewed brief helps it prepare and can replace a separate set of notes. The next three chapters examine the organizational change that follows. Chapter 8 commits capacity, Chapter 9 changes responsibilities and decisions, and Chapter 10 tests and adopts the new offering. These events are teaching material, not reported business results.

The organizational change at a glance

Organizational choice	Starting arrangement	Arrangement reached in Chapter 10
Customer involvement	Support follows an agreed service request	Customers can bring an unresolved expansion-planning question
Eli's contribution	Coordinate agreed delivery	Examine training implications before a request is finalized
Preparation	Individual notes and a duplicate summary	Reviewed brief plus a separate planning comparison
Resources and accountability	Planning and new material remain ideas within existing delivery responsibilities	Protected planning time, a limited design budget, named ownership, and a scheduled review

Choose What the New Capacity Is For

At the review, Maya brings a brief with her changes marked. Eli has confirmed the walkthrough status, and she has a question ready about the customer's possible expansion. The team has a preparation routine it can repeat, including a clear place for human review.

Lena agrees that account owners using the brief no longer need to produce a separate preparation summary covering the same information. Maya still reads the records needed to check the draft and chooses her agenda. Removing the duplicate summary lets the team put its attention into that review and the conversation it supports.

The change is not that preparation disappears. It is that the team stops asking a person to reconstruct the same account history in two places. The connected assignment gathers the permitted records, retains the source references, pauses for missing context, and delivers a draft for Maya's review. Maya spends her time on exceptions, interpretation, and the customer conversation.

Before the connected assignment	With the reviewed briefing routine	Capacity the team can choose to use
Maya searches several records and rewrites a preparation summary	The workflow reads the selected records, flags missing context, and prepares one source-referenced brief	Maya checks what matters, resolves exceptions, and prepares questions
Eli chases account owners when walkthrough status is unclear	The workflow asks the named owner, saves the answer, and resumes from the checkpoint	Eli reviews only the service questions that need his expertise
A second summary is copied into the preparation folder	The reviewed brief becomes the agreed preparation record	The duplicate step is retired; the team tests an earlier planning conversation

This table describes a proposed operating change, not measured savings. The team still records review effort and the time required to resolve exceptions.

Maya wants to use some of that attention to discuss customers' plans before they become urgent service requests. Eli has a related idea: improving the material that answers recurring training questions. Lena could ask them to keep those ideas on a list, but instead decides to give them time to investigate which would help customers most.

Be specific about the opportunity

The note about two possible additional locations gives Maya somewhere to begin. The customer has asked what support would be available, without yet deciding to expand. An earlier discussion could help it understand the training and service implications while there is still time to consider alternatives.

Maya proposes inviting the customer to a planning conversation with Eli. They would use approved account records to prepare, identify what they still need to learn, and discuss the customer's decision with the people responsible for it. The invitation would make clear that the conversation is exploratory and creates no commitment to an expansion date or service package.

The team should describe the opportunity at this level before calling it “higher-value work.” Maya can explain who the conversation is for and what decision it might help. Lena can then decide whether it deserves time alongside the team's existing obligations.

Prepare options from a customer's constraints

The assignment at a glance	
Inputs	Customer needs, attendance constraints, service options, availability, and unresolved dates
AI contribution	Develop feasible alternatives, explain competing priorities, and revise the comparison when conditions change
Human decision	Verify feasibility and decide which tradeoffs belong in the customer conversation
Evidence	Constraint violations, revisions, review effort, and contribution to the planning conversation

The planning service needs more than a blank worksheet. Maya and Eli want AI to assemble a customer-specific starting point from approved intake answers, service descriptions, and scheduling constraints. They can test that assignment with a small practice case:

- **C1:** If the expansion proceeds, each of two new locations needs one 90-minute remote training session.
- **C2:** The same customer coordinator must attend both sessions, so they cannot overlap.
- **C3:** For this exercise, the trainer has a two-hour window on Tuesday and Thursday mornings in each of the next two weeks.
- **C4:** The customer has not chosen opening dates. Both locations can attend either morning window.

They ask:

Prepare two scheduling approaches that satisfy C1–C4. Explain the tradeoff and list what must be confirmed before booking. Reference the constraints. These are planning options, not confirmed appointments.

An illustrative comparison is:

Approach	Proposed sequence	Reason to consider it	Confirmation needed
Same-week preparation	One 90-minute session Tuesday; the other Thursday	Both locations receive training within one week; the coordinator attends each	Which location goes first, actual dates, and customer confirmation
More time between sessions	One session Tuesday in the first week; the other Tuesday in the second week	The coordinator can apply lessons from the first session before the second	Whether a later second session fits the customer's plans

AI has combined the constraints into alternatives and explained how they differ. Eli checks duration and availability against the source records; Maya asks which tradeoff matters to the customer. A scheduling tool would still need to verify and reserve actual slots.

The simple schedule establishes whether the assignment can preserve stated constraints. A more useful test adds different kinds of information:

- **C5, customer note:** “Our senior coordinator has to attend both sessions, but she still covers the current location. If she is out twice in one week, nobody there can coach the rest of the staff.”
- **S1, standard service:** Two 90-minute sessions with no between-session review.
- **S2, staged service:** One session, a short learning note reviewed by Eli, and a second session the following week. The review requires one hour of Eli's time.

Maya asks AI to compare S1 and S2 against C1–C5, identify the coverage concern in the customer's own words, and show what the team must confirm. The illustrative comparison says S1 finishes sooner but places both absences in one week. S2 reduces that weekly coverage pressure and lets the second location benefit from the first session, but it takes two weeks and requires Eli's review time.

Then the customer adds a condition: the second location can attend only Thursdays, and the current location can release the coordinator only on Tuesday of the first week. The revised comparison schedules the first

location on that Tuesday and the second on Thursday of the following week. It explains why the same-week approach no longer satisfies the coverage condition and makes Eli's additional review commitment visible.

A template can hold fields, but it cannot by itself interpret the customer's explanation, connect it with two service designs, and revise the comparison when the condition changes. AI can prepare that first-pass synthesis repeatedly. Maya and Eli still verify the records and decide what to recommend.

Capacity has different forms

The team has removed duplicate preparation, but it has not established a number of hours saved or a financial return. Its immediate decision is how to use the capacity it can observe. The planning comparison lets Maya and Eli begin with alternatives tailored to the customer's constraints and test whether more of those conversations are practical.

The Accenture and Aspen Institute report calls a related skill “Re-humanizing time,” connecting technology-supported capacity with human interaction, creativity, and decisions. [*Advancing Missing Middle Skills for Human-AI Collaboration*](#), printed page 25.

Give people a say

Maya sees an opportunity in earlier customer conversations; Eli sees it in recurring questions that arrive after plans have been made. Lena agrees to begin with the conversation and use what they learn to choose a piece of training material. She reserves preparation and meeting time while keeping current commitments covered.

People should have a say in what happens to capacity they help create. They know where they have been rushing, what they have postponed, and where their experience could contribute more. Lena still weighs those ideas against the organization's priorities, but the team helps shape the choice.

Related application: Explain a capacity conflict

Lena needs to decide whether the team can support the new service alongside its existing commitments. For this exercise, the available planning time and recorded requests are:

Record	Available hours or required effort	Context
K1, Eli's available planning time	6 hours	After current delivery commitments
K2, request A	2 hours	Existing customer; training dates already agreed
K3, request B	3 hours	Proposed planning service; no agreed delivery date
K4, request C	2 hours	New customer needs an answer before choosing an opening date

The requests total seven hours against six available. A calculation identifies the one-hour shortfall. AI reads the surrounding context and asks whether request B can move, whether its scope can change, or whether someone with suitable expertise can help. Lena verifies the estimates and decides what to change. The contribution is an explanation tied to the commitments, not a bare utilization percentage or a judgment about anyone's performance.

Consider a different service

Maya's invitation changes the team's role. Instead of waiting for a confirmed expansion, it offers help while the customer can still change its plan. Customers need to understand what guidance the team can provide, what remains undecided, and who can authorize a commitment. Delivering that support may eventually require training-design expertise and a regular planning responsibility.

Make a capacity commitment

Lena records the decision to try the planning conversation. Maya will confirm the customer's interest and identify who should attend. Eli will prepare the service questions, and the team will use the reviewed brief as a starting point. They also agree to examine the preparation, meeting, and follow-up effort before deciding whether to offer it again.

The calendar now contains time for something that previously remained an idea. Lena has also removed the duplicate summary requirement, so the team has a concrete change to examine alongside the new activity. At the review, they can discuss whether that arrangement gave them enough room and what they would need for a second conversation.

Workbook: Choose what the new capacity is for.

Rethink Roles and Decisions

The planning conversation is on the calendar, and Maya wants Eli to join it. Until now, his involvement has usually followed an agreed service request. Asking him to help while a customer is still considering its plans changes what he needs to prepare and when his time must be available.

This is where I would put the org-chart conversation. The team has an activity it wants to carry out and can examine the responsibilities needed to support it. A new reporting line may eventually help, but the immediate decisions concern Eli's time, Maya's authority, and how the two will work together.

Describe the future activity

Before accepting the invitation, Eli asks what happens if an urgent service request arrives at the same time. He also wants to know whether Maya can ask him to prepare options or only describe support that has already been agreed. Lena responds with a short role agreement for the trial:

Role	Previous expectation	Trial responsibility	Authority and limit
Maya	Lead account conversations after service is defined	Frame the customer's planning question and lead the earlier conversation	May involve Eli within the reserved planning time; does not promise delivery
Eli	Coordinate agreed training delivery	Examine training implications and prepare options before the request is finalized	May explain options and dependencies; additional resources still require approval
Lena	Review current delivery commitments	Protect planning time, arrange coverage, and decide whether the activity continues	May change the allocation for the trial and approves any permanent role change

The new routine also needs context owners. Maya is responsible for the account facts and the customer's question; Eli owns the service and training assumptions he verifies; Lena owns the limits on available planning time. The system may retain their approved answers for the current run, but it does not decide when an old answer remains valid. Each owner can correct or withdraw the information in their area before the comparison is used.

The AI assignment prepares background and questions from the approved records. The role agreement determines what Maya, Eli, and Lena do with that preparation. If planning becomes a regular service, Lena will revisit the recurring allocation of Eli's time and how the contribution is recognized.

Related application: Turn a stalled opportunity into a specific next step

A pipeline rule flags a fictional opportunity after ten days without an update. That flag is ordinary automation. AI receives the opportunity record and permitted correspondence to help the owner understand what to do:

- **O1, CRM:** Proposal sent; no next meeting recorded.
- **O2, customer email:** “The operations lead needs to see the training approach before we can move forward.”
- **O3, internal note:** The training outline is ready for Maya's review.

AI identifies the operations lead's review as the stated dependency and connects it with the internal note. Maya's responsibility is no longer merely to notice a stale record; she verifies the dependency and decides whether to offer the relevant discussion. The CRM stores the approved next step and reminds its owner, while AI contributes the interpretation that the ten-day flag alone cannot provide.

Reconsider the decision path

A different kind of preparation can also make a revised decision process worth trying. Consider a second hypothetical example: a team preparing recommendations for an internal project committee. One person gathers material, another prepares a summary, and several reviewers add comments in sequence. By the time the committee meets, much of its attention goes to understanding the background.

The team could test an arrangement in which AI helps assemble an initial evidence brief and options from approved records. Relevant specialists would check their areas before committee members receive the reviewed material and identify questions for discussion. That proposal requires agreement on when records are ready, how specialists record their judgments, and what the committee needs before it meets. The following table sets out a possible arrangement:

Decision activity	Possible revised arrangement
Gather background	A configured AI-supported process prepares a draft from selected approved records, with source references.
Assess the evidence	Named specialists review the relevant facts and assumptions before circulation.
Identify unresolved questions	The decision owner collects the questions that require discussion rather than treating the draft as a recommendation already accepted.
Choose a course of action	The authorized people make and record the decision.
Learn from the result	The team revisits important assumptions and improves the preparation process.

The change to examine is whether the committee can spend more attention weighing options and less reconstructing the background. The team would need to observe that shift in its own decision process; a shorter briefing alone would not establish it.

Related application: Compare proposals against the actual decision

For the internal committee, supply AI with three fictional records:

- **Q1, proposal A:** Six-week implementation; training is outside the quoted scope.
- **Q2, proposal B:** Eight-week implementation; two training sessions included.
- **Q3, internal requirement:** Desired launch in seven weeks; staff need training before launch.

AI compares both proposals with Q3. It finds that A fits the desired implementation window but leaves training timing and cost unresolved. B includes training but runs one week beyond the desired launch. The decision owner can ask A for a training proposal and B whether its sequence can change. Specialists verify the interpretation and bring the unresolved tradeoffs into the meeting. AI aligns the separate documents with a common requirement; it does not invent prices or choose a winner.

Put authority where it is needed

Maya may arrange the planning conversation and involve Eli within the agreed time, but neither promises delivery during it. The specialist in the committee example has a comparable right to return a brief for revision and mark an assumption unresolved.

I would be reluctant to give someone broader responsibility while leaving them to negotiate permission at every step. The trial agreement identifies which decisions belong with the person who has the context and which still require a leader or specialist.

Broader perspectives in product development

In the 2026 published study *The Cybernetic Teammate: A Field Experiment on Generative AI and Teamwork*, Fabrizio Dell'Acqua and colleagues report a preregistered field experiment with 791 P&G professionals working on product-innovation challenges. Without AI, technical and commercial participants tended toward proposals reflecting their own specialties. AI-assisted participants produced proposals that more evenly combined those perspectives. The result supports examining how AI can broaden the considerations people bring to a decision; it is not evidence that every AI-assisted team or service will perform better. **Organization Science**, 2026.

Recognize the new responsibilities

Eli now helps a customer understand its choices before there is a service request to coordinate. Maya guides a planning conversation that may end with several options still open. Framing assignments, reviewing evidence, and helping colleagues learn may also become recognized responsibilities as the practice develops.

If the service continues, Lena will consider training for Maya and Eli and whether the team needs additional design expertise. That decision follows what customers need and what the trial reveals.

Check whether the organization still rewards the old routine

If Eli is recognized only through completed service requests, planning looks like time away from his main responsibility. Lena therefore includes it in the trial expectations and reviews whether the customer understood the options and next steps, along with the effort required. Committee specialists likewise need time to review assumptions before the meeting.

Test the arrangement before expanding it

Before the conversation, Maya and Eli identify what they can explore together and what requires another specialist. The trial tests that arrangement as well as the meeting itself. Its review will show whether the responsibilities need more time, support, or a permanent place in the organization.

Workbook: Rethink roles and decisions.

Create New Possibilities

In the planning conversation, the customer explains that it is still deciding whether to add the two locations. Its immediate concern is how existing employees would help train new colleagues without leaving the current locations short of support. It wants to understand that dependency before settling the expansion schedule.

Maya's original brief helped her remember to ask about expansion. With Eli involved earlier, the team can now explore what that expansion would require. Their preparation becomes part of a service that helps the customer make a decision while there is still room to change its plans.

Before designing every detail, Lena records the offering the team will test:

Early planning support	Trial decision
For whom	Existing customers considering additional locations
Customer need	Understand training, coverage, and sequencing before fixing an expansion schedule
AI contribution	Prepare a comparison from the intake, account records, and service descriptions
Human contribution	Maya frames the decision; Eli checks service and training implications
Resources	Protected planning time and limited training-design help
Review	Examine customer usefulness, preparation effort, corrections, and support required before wider adoption

Ask who could benefit

Eli asks which employees would support training and what responsibilities would still need coverage at their existing locations. The customer has not worked out all the answers. They agree to gather that information and use it to compare possible approaches before discussing a delivery commitment.

After the call, Maya and Eli propose an individualized planning comparison followed by a conversation. The customer would supply its locations, training needs, available people, and unresolved decisions. AI would use that intake with the approved service descriptions to draft options, explain dependencies, and prepare questions about missing information. Maya and Eli would check the comparison before discussing it with the customer.

The anchor example in Chapter 8 shows the difference between providing an empty form and preparing alternatives from the answers. The worksheet records the customer's situation; the AI assignment develops a first-pass comparison from it. That contribution needs to remain visible as the service grows, so the team can establish whether preparing and reviewing those comparisons is practical for more customers.

Broaden access to useful support

An internal learning team could use AI-supported rehearsal between facilitated sessions. Participants practice with a configured simulation, receive feedback framed by the learning objectives, and bring questions back to a facilitator. The team assesses the exercises, sets the information boundaries, and uses the next session to address what participants found difficult.

Related application: Adapt the same procedure for different roles

Give AI procedures P1–P3 from Chapter 3 and ask for two items: a coordinator's job aid for changing a remote-session date, and a manager's practice question about a change from remote to on-site delivery. Both must preserve the procedure and identify the source.

The coordinator aid covers availability, written confirmation, the action-list update, and the retained reference. The manager's practice question adds a change from remote to on-site delivery, which requires service-lead approval under P2. A learning designer checks both versions for accuracy and fit. AI adapts the explanation and exercise to the decisions each role makes while preserving the procedure. The team still needs to test accessible formats and other support with intended users.

Explore more than the first plausible option

AI can help compare facilitator-led practice, peer sessions, AI-supported rehearsal, and combinations of those approaches. Maya and Eli can likewise compare a group session for shared questions with an individual service for customer-specific plans.

Lena chooses the individual comparison because its value depends on each customer's circumstances. A group session remains an option for general training questions. The comparison identifies assumptions to test; it does not replace the decision about which experience the team can support.

Service option	Strength	Tradeoff the trial must expose
Individual planning conversation	Applies the customer's own constraints and unresolved decisions	Requires intake, preparation, specialist review, and a follow-up route
Group session	Reaches several customers with shared guidance	Cannot resolve account-specific commitments or coverage constraints
Self-guided resource	Available without scheduling a meeting	Needs accessible design, maintenance, and a clear path to a person when the resource is insufficient

The team begins with the individual conversation because the customer is making a specific planning decision. It keeps the other formats available for needs that do not require the same level of context.

Design the whole experience

The customer experience begins with an invitation and intake, continues through a reviewed comparison and conversation, and ends with recorded questions and next steps. Behind it sit the AI assignment, Eli's review time, separate account records, and a defined follow-up route.

Lena also budgets for help designing the training-planning material, so Eli is not expected to become an instructional designer because he understands service delivery. The learning example needs a comparable route back to a facilitator when participants cannot interpret the feedback.

Try a small version of the new offering

The trial covers the intake, AI-prepared options, review, conversation, and follow-up. One customer cannot complete the proposed-date field because an earlier decision remains open, so Maya and Eli revise the intake to capture that decision and its owner. They also find that Eli's review must happen before the meeting, and Lena adds that time to the arrangement.

The review does not establish a financial return or make the offering ready for every customer. It tells Lena what the team can support consistently and what needs another test.

Carry a repeated customer question into a new resource

The assignment at a glance	
Inputs	Permitted, de-identified observations from several planning conversations
AI contribution	Group related needs, cite the observations, and propose candidate resources
Human decision	Validate the pattern with customers and decide whether to invest
Evidence	Grouping corrections, customer response, development effort, and later use

The planning conversations may also reveal a need shared across customers. Suppose the team reviews these fictional, de-identified observations with permission:

- **F1:** Asked who should train new location coordinators.
- **F2:** Asked whether experienced employees could support training while covering their existing locations.
- **F3:** Asked for a way to record which decisions must precede training dates.
- **F4:** Asked who owns the training plan after the initial discussion.

The team uses this instruction:

Group the observations by the customer need they describe. For each group, cite the supporting observation IDs, propose one practical resource, and state what we still need to validate with customers. Do not treat four observations as evidence of market-wide demand.

An illustrative result is:

Customer need	Supporting observations	Candidate resource	Validate next
Assign training responsibility	F1, F4	Responsibility checklist for the customer and service team	Whether the checklist clarifies ownership after the planning call
Protect current-location coverage	F2	Coverage conversation guide	Which constraints recur and who needs to participate
Sequence decisions before dates	F3	One-page planning sequence	Whether the order matches how customers actually decide

The sample contains four observations, not evidence of market-wide demand. Maya checks the grouping and asks customers whether the proposed resources would help. A learning designer develops the approved resource, with AI assisting on draft explanations and examples.

The contribution has moved beyond an individual document: AI helps the team compare qualitative records across conversations and turn a repeated need into a candidate resource. The organization still decides whether to invest and checks whether customers use what it produces.

Decide which possibilities to pursue

After reviewing the trial, Lena makes the decision visible:

Decision	Adopted arrangement
Offering	Early training-planning support for existing customers considering additional locations
Ownership	Maya owns the customer experience; Eli's planning and review time is part of his agreed role
Resources	Limited training-design budget and protected development time
AI assignments	Customer brief, planning comparison, and approved learning-material drafts, each with a named reviewer
Next review	Reconsider eligibility, effort, customer use, and support before wider expansion

The organization now offers help at a point where it previously waited for a service request. It has allocated people and resources to that offering, changed expectations for existing roles, and created a reason to develop additional expertise. Those are the changes that make this example one of transformation. The AI-prepared brief was the starting point for a series of decisions about how the organization could serve its customers.

A leader should be able to explain those decisions as clearly as the technology involved. Customers can now bring an unresolved planning question to people who have time to help them examine it.

Workbook: Create new possibilities.

Closing · Bring it to the team

Begin with one meeting and one activity. Ask the person who does the work to bring a recent example. Include the colleague who receives the handoff and the leader who can change the surrounding routine.

By the end of that meeting, decide what AI should contribute, which records it may use, who will review the result, and when the group will examine what happened. Name one existing step that may be changed or retired if the trial works. That last decision gives the experiment somewhere to go.

The team may keep the assignment small. It may also discover, as Maya, Eli, and Lena did, that better preparation makes an earlier customer conversation practical. The next move should follow what people learn from the work in front of them.

Workbook · Implementation, Enablement, and Transformation

Use these pages with the people involved in the activity. Put a date and an owner on each page you use, and preserve the completed version with the instructions, examples, and run records it refers to. Record a changed answer instead of quietly replacing the earlier agreement.

You do not need every page. Sections 1–3 establish a first assignment. Sections 4–7 support work that other people rely on or that continues across steps. Sections 8–10 are for a choice about capacity, roles, decisions, or a new offering.

Completed first page · Maya's customer brief

This page shows the minimum needed for a useful first conversation. Later workbook sections add connected responsibilities, evidence, capacity, and organizational decisions when they become relevant.

First conversation	Maya's example
Activity	Prepare for a customer check-in
Person who understands it	Maya, the account owner
What she wants to do better	Arrive with open commitments and useful questions already visible
AI contribution	Prepare a source-referenced draft from selected notes and the open-action list
Human decision	Maya checks the draft and chooses the agenda
First evidence	Corrections, missing items, review effort, and whether the brief supported preparation
Trial decision	Try five selected check-ins, then revise, continue, or stop

Maya's page identifies the assignment, its human owner, and the evidence needed for a first decision. It does not claim time saved, revenue, or a general capability the trial has not established.

Implementation

1. Choose a useful starting point

Use this page to choose one recognizable activity. Describe what happens now before deciding what AI should do.

Date and page owner

Activity and the person who understands it

What people want to do better or more fully

Proposed contribution

Before or when _____, AI will help _____ by _____, using _____.

Why this is a useful first example

1. Choose a useful starting point · boundaries and decision

Existing permission and approved environment for trying it

Current handoffs, repeated preparation, or missing context worth observing

What is outside this first trial

Decision: proceed, revise the starting point, or stop

2. Give AI an assignment

Complete the agreement before connecting tools or permitting an action.
If a row is unresolved, name the person who will resolve it.

Date, assignment owner, and current version

Decision	Our agreement
Purpose and intended user	
Person or event that starts the assignment	
Approved inputs and who selects them	
Expected output or action, with an example	
Permitted actions and human approval points	
Named human owner and coverage	
Missing information or exceptions that return to a person	
Where the result goes and who checks it	
Quality requirements and how they can be checked	

One sentence the intended user could use to explain the assignment

The most important way this assignment could be misunderstood

3. Build the first working example

Use one bounded example whose facts a reviewer can verify. Keep the source records with this page or identify where they are stored.

Date, builder, and reviewer

Representative fictional or approved records

Instructions and configuration used

What happened when we tried the assignment

What the reviewer could use and what needed revision

A different case we should try

Decision about a wider trial, including participants and review date

3. Build the first working example · procedure coverage

Use this page to test both a question the available material can answer and a question it cannot. A useful system should answer the covered question and stop on the uncovered one.

Test	Source support expected	Result observed	Pass, revise, or stop
Covered case			
Uncovered case	No sufficient source		
Conflicting sources, if relevant			

What the answer must cite or show

Exact response expected when support is missing

Correction made to the instruction, source set, or assignment boundary

Result of the repeated test

Enablement

4. Decide what can change

Use this page with the people doing, receiving, managing, and supporting the work. Permission to participate is not the same as permission to change every surrounding process.

Date, trial owner, and people represented

Decision	Our agreement
Changes participants may make during the trial	
Changes requiring additional approval, and who can give it	
Time and resources for practice and review	
Person who can revise or pause the assignment	
How feedback will reach someone who can act on it	
What the team hopes to accomplish with the support	
People who should help define data, access, customer, employment, or legal limits	

4. Decide what can change · pause and participation

Existing step that may be changed or retired if the trial works

Condition that would pause the trial immediately

How a participant can raise a concern without becoming the permanent support owner

What participants receive when the assignment is paused

Person who diagnoses the issue and decides whether to restart

Where the concern, decision, and resulting change will be recorded

5. Help people work well with AI

Plan a practice session around a real example and a difficult moment.
The aim is judgment: what to accept, what to check, and when to stop.

Date, practice owner, and intended participants

What a participant should be able to do

Example we will demonstrate and let people try

A question or uncertain result participants should practice handling

Support available for different starting points

5. Help people work well with AI · support and learning

Named support owner, time for that responsibility, and backup

Next practice opportunity

How people can propose, test, or improve the assignment without becoming its permanent maintainer

What the facilitator will observe during practice

Change to make before the next practice session

6. Connect the responsibilities

Trace an example through the entire activity, including what happens after the first output.

Date, workflow owner, and run or example ID

Step	Person or system acting	Input	Output and recipient	What confirms readiness for the next step?
1				
2				
3				
4				
5				

Trigger and completion event

Where the run state is saved if the work pauses

6. Connect the responsibilities · completed movement map

This completed map condenses Maya's briefing run. It shows the handoffs that matter before a team draws its own version.

Step	Actor and kind of work	Evidence or state	Next movement
1	Scheduler, ordinary software	Meeting ID and assignment version	Start one run
2	Connectors, ordinary software	Selected CRM notes and open actions	Return retrieved records
3	AI, judgment	Missing walkthrough status identified	Ask Eli; pause
4	Eli, human decision	Approved answer stored with source and time	Resume from checkpoint
5	AI plus evaluator, judgment	Draft, cited sources, correction request	Revise or stop at limit
6	Maya, human decision	Reviewed brief and approval	Prepare final write
7	Connector, ordinary software	Idempotency key and destination check	Write once; verify receipt

What the map makes visible that a list of tools would not

The run has two human decisions, a saved pause, an evaluation limit, and a verifiable completion event. A diagram of connected products would not by itself show any of those responsibilities.

Question this completed example raises for our activity

WORKBOOK · CONTINUED

6. Connect the responsibilities · our movement map

Draw or write the path from trigger to confirmed completion. Include branches for missing information, failed checks, and a declined action.

Movement	Actor	Input or saved state	Output, recipient, or branch
Trigger			
Retrieve or receive			
Interpret or draft			
Check			
Ask or pause			
Resume			
Approve or reject			
Act			
Verify completion			

6. Connect the responsibilities · decision points

Where ordinary software ends and AI judgment begins

Where a human decision is required

6. Connect the responsibilities · interrupted run diagnosis

Use a real failed or interrupted run when one is available. Until then, rehearse a plausible one. Diagnose the point of failure before changing the whole assignment.

Question	Our finding
What was the run trying to complete?	
Last confirmed step	
Saved state and instruction version	
Missing, conflicting, or failed input	
Did any external action occur? What proves it?	
Person responsible for the next decision	
Safe next step	

Classify the outcome

Complete: all required checks and the final receipt passed.

Stopped: the system followed a rule that required it not to continue.

Failed: the run could not complete as intended and needs diagnosis or recovery.

Our classification and evidence

6. Connect the responsibilities · stop, resume, and uncertain action

Agree on recovery before the first consequential write. Resuming should continue the same run from a trustworthy checkpoint, not silently begin a different assignment.

Recovery decision	Our agreement
Conditions that stop the run	
State that must be saved before a pause	
Person allowed to approve, reject, or revise	
Information that may be added on resume	
Conditions that require a fresh run instead	
Maximum revisions, tool calls, or elapsed time	

If a write times out, what record will we inspect before retrying?

Idempotency key or other duplicate-prevention method

Evidence that confirms the action happened exactly once

6. Connect the responsibilities · authority and handoff

Responsibility level	Current choice	What must be true before AI takes another step
Draft		
Recommend		
Prepare an action		
Act within a boundary		

Where a question goes, who answers it, and where the answer is recorded

How a draft, a reviewed result, and an agreed action are distinguished

7. Build a practice that improves

Start by defining completion. A result that looks plausible but lacks a required check is not accepted.

Date, evaluation owner, and assignment version

Completion requirement	How it is checked	Evidence retained	Owner
Required records retrieved			
Factual statements supported			
Required sections present			
Human review completed			
Final action verified			

Failure or stop conditions

Maximum evaluator-and-revision cycles

7. Build a practice that improves · trial record

Use consistent definitions when comparing approaches. Include review effort and keep setup, training, and support effort visible.

Run or approved reference	Context and assignment version	Outcome, checks, and review effort	Correction or next adjustment

Run IDs and location of complete traces

What we observed and what we have not established

What each run should retain so someone can reconstruct the result and correction

7. Build a practice that improves · decision and maintenance

Decision and maintenance

Decision, owner, and next review date

Location of current instructions and representative examples

Changes that should trigger another check

How a new participant will learn the process

7. Build a practice that improves · accepted-result cost

Count the cost of every attempted run, including retries and human review, then divide by the number of accepted results. Keep estimates labeled as estimates.

Cost component	Quantity	Rate or unit cost	Total
Model and tool use			
Reviewer time			
Correction and recovery			
Other run-specific effort			
Total attempted-run cost			

7. Build a practice that improves · result counts and decision

Result count	Number
Attempted runs	
Accepted results	
Stopped runs	
Failed runs	

Accepted-result cost: total attempted-run cost ÷ accepted results

What the number includes and excludes

Decision this evidence can support

Transformation

8. Choose what the new capacity is for

Use observed changes, not assumed savings. Name the work that changed and the opportunity that now deserves a bounded test.

Date, decision owner, and people involved

Benefit we have observed

What people would like to give more attention to

One opportunity we will explore

Who would benefit and whose agreement we need

8. Choose what the new capacity is for · commitment and limits

Owner, time, resources, and next decision date

Existing step we will remove, reduce, or keep

Exceptions and context work that still require people

What we are not claiming about capacity, cost, or return

9. Rethink roles and decisions

Describe the future activity before proposing a title or reporting-line change. Include the time, authority, and expertise the activity needs.

Date, activity owner, and people affected

Future activity or decision we want to support

Responsibility or decision	Current arrangement	Proposed arrangement	Authority, time, or expertise required

9. Rethink roles and decisions · participation and evidence

What AI would contribute and which judgments people would make

Role descriptions, schedules, or expectations that need to change

How the people affected will help shape the arrangement

Bounded example for trying the revised arrangement

What would tell us whether the change helped

9. Rethink roles and decisions · authority map

Use one real decision. Separate preparation from recommendation, approval, action, and verification.

Decision or action	AI may prepare	AI may perform within a boundary	Human approval required	Named owner

Person who owns the source context and can correct it

Person who can return the work for revision or stop it

Coverage when the named decision owner is unavailable

Role, schedule, recognition, or support change required for a fair trial

10. Create new possibilities

Define one small, complete experience. The test should include intake, preparation, review, delivery, support, and follow-up—not only an AI-produced artifact.

Date, offering owner, and intended participant

Person or group whose need we want to understand

What we have learned from them and what remains an assumption

Proposed service, experience, or internal capability

10. Create new possibilities · delivery and evidence

Alternative approaches worth considering

Human and AI contributions to the proposed experience

Requirements for access, support, participation, and appropriate use

Small version we can test with permission

What we will observe, including the effort to deliver it

10. Create new possibilities · service decision

Compare at least two ways to meet the need before choosing what to test.

Option	Intended participant and need	Human and AI contributions	Access, support, and main tradeoff

Chosen small version and why

Assumption the test is designed to examine

10. Create new possibilities · evidence and decision

Evidence and feedback required before adoption

People and resources required for responsible delivery

Decision: adopt, revise, repeat, or stop

Decision owner, next action, and review date

Application index • What AI contributes

AI does not need to perform every step for a process to be AI-supported. A calendar integration finds an available slot, a database stores a record, and a calculation totals approved hours. AI contributes where the assignment requires interpretation, retrieval, comparison, drafting, adaptation, or movement across several permitted steps.

AI contribution	Application	Go to	Page
Interpret an open-ended request	Customer check-in preparation	Chapter 2 • Prepare the customer check-in brief	15
Retrieve and apply source material	Procedure questions	Chapter 3 • Procedure check	21
Compare differently worded records	Data consolidation, billing, and proposal review	Chapter 1 • Find the repeated preparation	9
Generate a constrained draft	Customer check-in brief	Chapter 3 • Correct the brief	19
Practice, explore, and adapt	Conversation rehearsal and role-specific learning	Chapter 5 • Practice with an example	29
Carry work across several steps	Scheduled briefing with pause, review, and verified write	Chapter 6 • The successful run	44
Diagnose an interrupted run	Stop, resume, and uncertain-write recovery	Chapter 6 • The interrupted run	45

AI contribution	Application	Go to	Page
Evaluate and revise	Completion checks and accepted-result cost	<u>Chapter 7 · Define completion</u>	49
Develop constrained alternatives	Customer training schedule	<u>Chapter 8 · Prepare options</u>	61
Explain a capacity conflict	Planning requests against available time	<u>Chapter 8 · Give people a say</u>	63
Connect records to a next step	Stalled opportunity	<u>Chapter 9 · Describe the future activity</u>	66
Adapt an explanation	Procedure job aid and practice question	<u>Chapter 10 · Broaden access</u>	73
Group qualitative evidence	Candidate customer resources	<u>Chapter 10 · Carry a repeated question</u>	76

The page number takes the reader to the beginning of the named section. The link goes to that destination in the digital edition.

Working terms

- **Agent:** Software configured to use AI and available tools to carry out steps within a defined assignment.
- **Accepted result:** An output or action that passes the checks agreed for its use, including any required human review.
- **Baseline:** A description or measure of the current activity used for comparison with the AI-supported version.
- **Checkpoint:** A saved snapshot of a run's state that allows the same run to continue after a pause or failure. A checkpoint is not automatically durable memory for other assignments.
- **Connector:** A configured route that lets one system retrieve information from or take an action in another.
- **Context:** The records, instructions, prior messages, tool results, and other information available for the current judgment. Context may be selected differently for each run.
- **Evaluation:** A check of a run or result against explicit requirements. An evaluator may identify a revision, but the assignment still needs a stopping rule and any required human review.
- **Human approval:** A named person's decision to permit, reject, or revise a proposed action. The approval should apply to a particular state and action, not to an unspecified future change.
- **Idempotent action:** An action designed so that a retry does not create a second unintended effect. An idempotency key and a post-action receipt are common safeguards.
- **Loop within a run:** Repeated movement during one attempt, such as requesting a tool, receiving the result, evaluating a draft, and revising it until a completion or stop condition is reached.
- **Model training:** A separate development process that changes or creates a model from training data. Saving a run, instruction, or answer does not train the model.

- **Reusable instruction:** Versioned guidance used across runs. It may tell the system what to do, but it is distinct from the live context, saved checkpoint, and model.
- **Run record:** The source references, instruction and configuration versions, result, review, and action needed to reconstruct a particular attempt.
- **Source record:** The note, document, or system entry used to support a factual statement.
- **Stored memory:** Application-managed information retained for later runs or users according to a defined rule. It should have an owner, scope, and way to correct or remove it.
- **Tool:** A function available to the configured system, such as retrieving records, checking a calendar, or preparing a write. Tool availability does not itself grant authorization.
- **Workflow:** A defined arrangement of triggers, steps, handoffs, and responses. Some workflows follow fixed paths; others contain model-directed choices.
- **Workflow graph:** A representation of the steps and possible routes through a workflow, including branches, parallel paths, pauses, and return paths. A graph can contain ordinary software, AI judgment, and human decisions.

Acknowledgments and influences

Geoff Woods' *The AI-Driven Leader* has had a substantial influence on how I think about AI, particularly the idea of using it as a strategic thought partner. Ethan Mollick's *Co-Intelligence* has also helped shape my thinking about people working and learning alongside AI. I am grateful for the questions both books have encouraged me to pursue.

My emphasis in this book is on the organizational decisions that accompany that collaboration: giving AI a recurring place in an activity, helping people work well with it, and considering what the organization could do differently. Woods and Mollick address related territory in their own work, and I encourage readers to explore both books directly. The reading notes below identify those influences and other sources cited in the chapters.

Sources and further reading

These sources identify influences and related thinking. They are not evidence that the hypothetical examples produced real-world results. The chapter sequence and workbook are practical proposals to adapt and test in context.

Changing technical documentation was checked on September 6, 2026. Product behavior and availability may change; use the linked documentation to confirm the current implementation details.

Books and practitioner perspectives

Geoff Woods. *The AI-Driven Leader*. Entrepreneur Press, 2025. Woods' emphasis on strategic thought partnership is a major acknowledged influence. His book also addresses organizational change. The authorized excerpt describes that scope; readers interested in his named frameworks and examples should consult the book directly. [Authorized excerpt, Porchlight Book Company, July 8, 2026.](#)

Ethan Mollick. *Co-Intelligence: Living and Working with AI*. Portfolio, 2024. An acknowledged influence on the book's human-AI collaboration perspective, including AI in work and learning. [Publisher page, April 2, 2024.](#)

Ethan Mollick. “Making AI Work: Leadership, Lab, and Crowd.” *One Useful Thing*, May 22, 2025. Related thinking on leadership, employee participation, and experimentation in organizational AI adoption. [Author essay.](#)

Accenture and the Aspen Institute. *Advancing Missing Middle Skills for Human-AI Collaboration*. 2018. The report connects its collaboration skills to Paul Daugherty and H. James Wilson's *Human + Machine*. Its discussion of “Re-humanizing time” is relevant to the capacity choices in Chapter 8. [Report, printed page 25.](#)

Alex Lieberman. “30 Features of an AI-Native Company.” August 31, 2026. A practitioner blueprint used in Chapters 6 and 7 to examine context stewardship, evaluation, graduated responsibility, accepted-result cost, and traceability. It is an operating perspective, not empirical evidence that every organization needs the same architecture or practices. [Author article.](#)

Empirical research

Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. “Generative AI at Work.” *The Quarterly Journal of Economics* 140, no. 2 (2025): 889–942. The published customer-support study discussed in Chapter 5. This manuscript uses the journal version rather than the earlier working-paper figures. [Journal article.](#)

Fabrizio Dell'Acqua and colleagues. “Navigating the Jagged Technological Frontier.” *Organization Science* (2026). The consulting-task experiment discussed in Chapter 7, including the importance of task-specific evaluation. [Journal article](#).

Fabrizio Dell'Acqua and colleagues. “The Cybernetic Teammate: A Field Experiment on Generative AI and Teamwork.” *Organization Science* (2026). The published P&G product-innovation experiment discussed in Chapter 9. [Journal article](#).

Engineering mechanisms

These entries explain possible construction patterns. They do not establish that a configured assignment is reliable, authorized, or useful in a particular organization.

OpenAI. “Running agents.” OpenAI Agents SDK documentation. Describes the model-call, tool-call, tool-result loop and configured turn limits used to explain a loop within a run in Chapter 6. [Documentation, accessed September 6, 2026](#).

OpenAI. “Tools.” OpenAI Agents SDK documentation. Describes hosted and function tools, including retrieval and runtime tool availability. Tool visibility is not a substitute for authorization. [Documentation, accessed September 6, 2026](#).

OpenAI. “Human-in-the-loop.” OpenAI Agents SDK documentation. Describes pausing a run for approval, serializing run state, and resuming after a decision. It does not imply that every connector can maintain a durable wait. [Documentation, accessed September 6, 2026](#).

Anthropic. “Building effective agents.” December 19, 2024. Technical background for predefined workflow paths, model-directed tool use, evaluator-optimizer patterns, checkpoints, and stopping conditions. The publisher notes that implementation details have changed since publication. [Engineering article](#).

Anthropic. “Demystifying evals for AI agents.” January 9, 2026. Describes multi-turn agent evaluation against explicit outcomes and environment state. A model grader is one signal, not independent proof. [Engineering article](#).

LangChain. “Workflows and agents.” LangGraph documentation. Describes conditional routing, parallel branches, and evaluator-optimizer loops used in the Chapter 6 capability map. [Documentation, accessed September 6, 2026](#).

LangChain. “Persistence.” LangGraph documentation. Distinguishes thread-scoped state snapshots from application-managed information stored across threads. [Documentation, accessed September 6, 2026](#).

LangChain. “Interrupts.” LangGraph documentation. Describes pauses, saved-thread resumption, and the need for idempotent side effects when execution may repeat. [Documentation, accessed September 6, 2026](#).

Microsoft. “RAG and Generative AI.” Azure AI Search documentation, updated August 4, 2026. Technical background for targeted retrieval, cited sources, and an agentic-retrieval option that can plan parallel subqueries. The pattern does not require this vendor, and setup and availability are product-specific. [Documentation](#), accessed September 6, 2026.

OpenAI. “Fine-tuning.” API reference. Describes the separate process of creating a model from training data through a fine-tuning job. It supports the distinction between model training and saving context, instructions, or a checkpoint. [Documentation](#), accessed September 6, 2026.

Keeping the capability map current

Use the mechanism sources as implementation references, not compatibility guarantees. Record the product or library version used for a trial, confirm that the documentation matches the configured runtime, and test the permissions, error paths, pauses, resumes, and receipts the assignment depends on. A feature appearing in documentation does not establish that it is enabled in a particular environment or dependable for a particular activity.

When a source or runtime changes materially, decide whether the assignment needs another representative-case check. Preserve the earlier instruction version and run records so the team can distinguish a software change from a change to the business agreement. The evidence for continued use should come from the configured assignment and its observed results, not from the documentation alone.