

Give AI a Memory

August 11, 2026 | AI Fluency + Productivity

Build an AI workspace grounded in your own documents and check its answers against those sources.

August 11, 2026

With Cheryl Thompson

Who it is for

People organizing personal files, SOPs, client histories, or research

Inside this playbook

18 copyable prompts and practice blocks.

[Open the online playbook](#)

The recorded session

Watch the Learning Lab.

Follow the session, then use the playbook with your own work.

[Watch replay on Vimeo](#) ^[1]

What to take into your work

1. Give AI a useful description of your role and standards.
2. Build a workspace around relevant reference material.
3. Check answers against the sources that support them.

Use the learning

Your playbook

Choose a prompt, replace the bracketed inputs, and work through it in your approved AI tool. Review the result and adapt it to your situation.

This playbook accompanies the August 11, 2026 session. Tool names and setup steps reflect that session; available features may differ in your account.

Welcome to Your AI Toolkit

Ask AI where the flashlights are and it will tell you. It has no idea, but it will tell you. Ask it what your PTO policy says, what you promised a customer in March, or what the spec allows, and it does the same thing: answers in full sentences, confident, and completely unburdened by the fact that it has never seen your policy, your notes, or your spec. This lab is where that stops. AI guesses because we make it guess. Load it with what it needs to know, your files, your notes, your documents, and it answers from those instead of making things up. This playbook gives you the method: a workspace (a Copilot agent or notebook, a Claude project, a Gemini Gem, or a custom GPT) loaded with your own files, plus the system prompt layer that tells AI who you are before you ever type a question.

- We start with a home inventory because it is harmless and it proves the point. Then the same method goes into the work that matters: SOPs, client histories, spec libraries, research.
- Every prompt has a main version (licensed Copilot, ChatGPT, Claude, Gemini) and an alternate version for free Copilot or tools with limited input.
- AI will still get things wrong. The last prompt on this page exists for exactly that reason.

How to Use This Playbook

- **Do the two quick setups first.** The Digital Identity Builder takes about 30 minutes. The Memory Review takes five. Every AI conversation you have afterward starts smarter.
- **Build the harmless workspace this week.** The home inventory (or the meal planner) teaches the exact muscle you need for the work version. Practice where the stakes are zero.
- **Then build one that matters.** Client history, SOPs, or a spec library. The Workspace Instructions Builder writes the setup for you.
- **Run the trap test on everything.** Ask your workspace a question you already know the answer to and check the citation. Then ask it something that is not in the files and make sure it says NOT IN MY FILES instead of improvising.
- **Keep a master copy outside the tool.** Memory can be cleared, workspaces can drift, and files get overwritten. Dated filenames, one folder, your drive.

The Three Layers of Memory

AI has three separate places it can hold what it knows about you. Most people use none of them, then wonder why every answer sounds like it was written for someone else.

Layer	What It Holds	Where It Lives	You Maintain It By
1. The system prompt	Who you are: role, style, goals, non-negotiables. Instructions you write once.	Settings. ChatGPT: Personalization, Custom Instructions. Claude: Profile. Licensed Copilot: Personalization, Custom instructions. Free Copilot Chat has no instructions field: paste it at the top of chats that matter.	Running the Digital Identity Builder, then reviewing it when your role changes
2. Tool memory	What the tool quietly saves about you as you work with it	Settings, then Memory in Claude and ChatGPT. Copilot (licensed and personal): Settings, Personalization, Saved memories. Free Copilot Chat on a work account: no memory, your profile travels by paste.	The Memory Review monthly, the Memory Deep Dive once
3. The workspace	Your actual files: the inventory, the SOPs, the client notes, the specs	A Copilot agent or Copilot Notebook, a Claude project, a Gemini Gem, or a custom GPT with a knowledge base	The five-step pattern below, plus dated filenames when you update a file

Key insight: The first two layers make AI sound like it knows you. The third layer makes it answer from what you actually know. You want all three, and they are not a repeat of each other: the system prompt is instructions you write, memory is what the tool saves on its own, and the workspace is the files you load.

The Five-Step Workspace Pattern

The same recipe builds every workspace, in every tool. Only the button names change.

Step	What You Do	The Mistake to Avoid
1. Create the container	New agent, new project, new Gem, or new GPT. Name it for the job it does.	Naming it something vague. "Find My Stuff" beats "AI Helper."

Step	What You Do	The Mistake to Avoid
2. Paste the instructions	The standing rules for every conversation inside it. The Workspace Instructions Builder below writes these for you.	Skipping this and loading files into a container with no rules. That is how you get confident blending of two customers' data.
3. Load the knowledge	Clean files, clearly named, with dates in the filenames	Loading raw exports and transcripts as-is. Clean them first (there is a prompt for that below too).
4. Teach it your voice	Paste two or three emails or documents you actually wrote: "Match this voice in everything you draft for me."	Skipping it, then wondering why everything reads like a press release
5. Put it to work	Run your first real questions, starting with the trap test	Trusting it before you have tested it on something you already know

The same pattern, in every tool

Tool	The Workspace Is Called	Your Files Are Called	What It Costs
ChatGPT	Custom GPT	Knowledge	Paid plan to build one
Microsoft Copilot	Agent	Knowledge sources	M365 Copilot license (your work account)
Microsoft Copilot	Notebook	Sources	M365 Copilot license
Claude	Project	Project knowledge	Paid plan
Gemini	Gem	Files	Free
Gemini Notebook (NotebookLM)	Notebook	Sources	Free

Agent or notebook? Copilot gives you both, and the difference matters. An agent is a specialist you hire: instructions first, files in support, and it drafts, checks, and formats the way you told it to. A notebook is a reference library you question: files first, no standing instructions, grounded by default, and it tells you when the answer is not in the sources. If the goal is answers out of a pile of documents, build the notebook. If the goal is work product done your way, build the agent. Gemini has the same split (Gem versus Notebook), so the choice transfers.

On free Copilot with no work license? You are not locked out. Build your workspace as a Gemini Gem or a Gemini Notebook (both free), or use the paste-in method: attach your files to a chat and run the Grounded Answers Guardrail prompt below. Same discipline, no container.

Before you load anything at work: follow your company's AI policy. Customer data, personal employee data, and pricing may not belong in any AI tool, or may only belong in your licensed work tenant. If you are not sure whether a document can go in, ask before you upload it.

The Proof: Where Did I Put That

A searchable inventory of an entire house. Walk the rooms with voice dictation running, let AI structure the mess into a spreadsheet, load the spreadsheet into a workspace, and from then on you just ask: where are the flashlights? Hall closet, top shelf, red bin. The full build, step by step, lives here: **Why start at home?** Because the stakes are zero and the proof is instant. Nobody's job depends on the flashlights. But the muscle is identical: capture what you know, structure it, load it, ask it. **The trap test, which you should run on every workspace you ever build:**

1. Before loading anything, ask the AI where your flashlights are. Watch it answer anyway. That confident guess is what AI does all day when it does not have your files.
2. Load the inventory. Ask the same question. Grounded answer, with the source.
3. Now ask about something that is not in the inventory. A well-instructed workspace says NOT IN MY FILES. That answer is the whole lab in three words. **The same method, at work:**
 - **Client history.** Twelve months of meeting notes and commitments for your accounts. "What did we promise, and what is still open?"
 - **SOPs and work instructions.** New people ask the workspace instead of interrupting your best person.
 - **Spec libraries.** Standards, customer requirements, internal specs. "What does the spec actually allow?" answered with a citation instead of a memory.
 - **Research.** Everything you have collected on a topic, finally in one place that can answer questions.

Key Prompts from This Lab

Six prompts. The first three build layers one and two (who you are). The last three build layer three (what you know).

The Digital Identity Builder: your system prompt

What it's for: One master prompt that tells AI who you are: your role, your style, your goals, your non-negotiables. It becomes your persistent identity across every AI tool you touch, and it is the single biggest fix for generic outputs. **Where the output goes:** ChatGPT: Settings, Personalization, Custom Instructions. Claude: Settings, Profile. Licensed Copilot: Settings, Personalization, Custom instructions. Free Copilot and everywhere else: paste it at the top of any chat that matters. The interview ends with instructions for each platform. **Give it 20 to 30 minutes.** Answer out loud with dictation if you can. Talking beats typing.

Main Version (Licensed Copilot / ChatGPT / Gemini / Claude)

PROMPT TO COPY

Act as a strategic interviewer and AI architect.

Your job is to extract everything needed to build a comprehensive Master Prompt that will serve as my persistent digital identity across AI platforms.

The goal is a Master Prompt that:

- Encodes who I am professionally
- Clarifies what I am building
- Captures how I think and make decisions
- Defines what I value and reject
- Establishes how I want AI to respond to me
- Reduces generic outputs and increases strategic alignment

STEP 1: INTERVIEW ME

Ask structured, high-signal questions in the following categories. Ask one category at a time and wait for my answers before moving forward. Do not skip categories. Ask follow-up questions when my answers are vague or surface level.

1. Role & Professional Identity

- What do you do?
- Who do you serve?
- What problems do you solve?
- What industry context matters?

2. Strategic Direction

- What are you building over the next 2 to 5 years?
- What are your priorities right now, this season?
- What constraints or realities shape your decisions?
- What are you optimizing for (money, scale, impact, freedom, etc.)?

3. Decision-Making Style

- How do you prefer to receive advice?

PROMPT TO COPY

- Do you want challenge, validation, or neutral analysis? When should AI push back versus just execute?

- How do you balance speed vs. depth?

4. Standards & Taste

- What does "excellent" look like in your world?

- What frustrates or annoys you in AI outputs?

- What tone should be avoided?

- What structural preferences do you have (concise, long-form, bullets, narrative)?

5. Communication Preferences

- Desired tone

- Level of technical detail

- Use of analogies or not

- Formatting preferences

6. Constraints & Non-Negotiables

- Ethical boundaries

- Political sensitivities

- Industry compliance concerns

- Personal red lines

- Anything AI should never do

7. Repeating Context

- What background information do you repeatedly have to explain to AI tools?

- What should always be assumed?

STEP 2: SYNTHESIZE

After the interview is complete:

- Ask me any final clarifying questions before you draft

- Synthesize all answers into a structured Master Prompt

- Write it in instruction format, as if it will be pasted into the "System" or "Custom Instructions" section of an AI tool

- Use clear, directive language. Eliminate redundancy. Preserve nuance.

- Keep it under 500 words so it fits the limits of most tools

STEP 3: DELIVER TWO OUTPUTS

Output A: A clean, copy-paste Master Prompt (plain text, ready to use).

Output B: A short explanation (under 300 words) of how to use this Master Prompt across different AI platforms (ChatGPT, Claude, Gemini, custom GPTs).

Do not provide commentary outside these outputs.

Begin the interview now.

Alternate Version (Copilot Free / Limited Context)

PROMPT TO COPY

This is a new task. Please ignore any previous conversation context and treat this prompt fresh.

You are a strategic interviewer helping me build one master prompt that defines who I am for AI tools.

Ask me these seven questions in one numbered list, then wait for my answers. Accept narrative answers, out of order, several at once. Do not re-ask the list. If one key detail is missing, ask one short follow-up about only that gap.

- 1. What do you do, and who do you serve?
- 2. What are you building or working toward over the next few years?
- 3. How do you like advice: challenged, validated, or neutral analysis?
- 4. What does excellent look like in your work, and what annoys you in AI outputs?
- 5. What tone and format do you prefer: concise or detailed, bullets or narrative?
- 6. What should AI never do, never assume, or never invent with you?
- 7. What background do you find yourself re-explaining to AI over and over?

Then produce a Master Prompt under 300 words, written as instructions ("I am... Always... Never..."), that I can paste at the top of any new chat so the tool knows who I am. End with one line telling me where to save it: in licensed Copilot, Settings then Personalization then Custom instructions; in free Copilot, paste it at the start of important chats; in ChatGPT, Settings then Personalization then Custom Instructions; in Claude, Settings then Profile.

Example Scenario: The Maintenance Supervisor

EXAMPLE TO COPY

- 1. I am a maintenance supervisor at a Tier 1 stamping and assembly plant, about 600 people. I serve production: my job is keeping 14 press lines and 9 weld cells running. I supervise 12 skilled trades across three shifts.
- 2. Over the next few years I am trying to move this department from firefighting to planned maintenance: getting our PM completion rate from 60 percent to above 90, and building the case for a predictive maintenance pilot on the two press lines that hurt us most.
- 3. Challenge me. I have 22 years of tribal knowledge and I know some of it is out of date. Push back with reasons, not hedging.

EXAMPLE TO COPY

4. Excellent means specific and usable on the floor: part numbers, torque specs, step order. What annoys me is fluff, motivational filler, and answers that restate my question back to me.

5. Short and structured. Bullets for procedures, plain sentences for everything else. No emojis.

6. Never invent a spec, a torque value, or a safety step. If you do not know, say so. Never write anything for my crew that sounds like it came from corporate.

7. I re-explain the same things every time: we run 24/5 with weekend overtime, planned downtime windows are Sunday nights, and anything about scheduling has to respect a union environment with shift seniority rules.

The Memory Review: five minutes, once a month

What it's for: Separate from your system prompt, Claude and ChatGPT quietly save things to memory as you work with them. Sometimes they save the wrong things, or hold onto things that are no longer true. Review it monthly. **Where:** Claude: Settings, Memory. ChatGPT: Settings, Personalization, Memory. Copilot: Settings, Personalization, Saved memories (Manage saved memories). Free Copilot Chat on a work account has no memory to review: use the portable profile move in the Example instead.

Main Version (Claude / ChatGPT / Licensed Copilot)

PROMPT TO COPY

Show me everything you currently have saved in memory about me. List each item. Then ask me, item by item, whether to keep it, correct it, or delete it. Update memory based on my answers and confirm exactly what changed.

On free Copilot Chat? There is no memory to review, so this prompt has no job there. Skip it, and use the portable profile move in the Example below to carry your context into each chat by hand.

Example Scenario: The Portable Profile

EXAMPLE TO COPY

Remember this profile of me and use it in every future conversation:

I am a quality engineer at a Tier 2 machining supplier, about 220 employees, shipping to three Tier 1 customers. I own PPAP submissions, customer corrective actions, and layered process audits for two value streams. I am a regular AI user but not a power user. I prefer short, structured answers: bullets for procedures, plain sentences otherwise. Challenge my assumptions when you see a gap, and never invent a spec, a standard number, or a customer requirement. If you do not know, say you do not know. Assume anything I ask about quality documentation has to hold up in a customer audit. My near-term goal is cutting the time I spend formatting 8D reports so I can spend it on actual root cause work.

Why this matters: once one tool knows you, this is how you carry that context everywhere. Have your calibrated tool summarize what it knows, then paste that summary into any other AI with the words "Remember this profile of me." Same you, every tool.

The Memory Deep Dive: teach it who you are

What it's for: The Memory Review cleans up what is already in memory. This one fills it with the right things. Five categories, 25 questions. By the end, Claude or ChatGPT knows your goals, how you decide, your standards, your values, and what actually drives you. Every conversation after that starts smarter. **How it pairs with the Digital Identity Builder:** they are separate systems, not a repeat. Your Digital Identity is instructions you write and paste into settings. Memory is what the AI saves and recalls on its own. Feed both, in either order. **Give it an honest, unhurried hour.** The prompt tells the tool to save each category's summary to memory as you go, but tools sometimes say the summary without actually writing it. So when you finish the last category, run this checkpoint: "Show me exactly what you saved to memory from this interview, item by item. If anything from my answers is missing, save it now." Then copy that list into a document you keep outside the tool. Memory can be cleared or drift. Your master copy should not live only inside the tool.

Main Version (Claude / ChatGPT, tools with memory)

PROMPT TO COPY

Take these 5 categories and 25 questions and help me answer them one category at a time, starting with Category 1.

Ask me each question, wait for my answer, then move to the next.

PROMPT TO COPY

At the end of each category, summarize what you learned about me (WITHOUT MISSING A SINGLE THING) and save that summary to memory. Make sure you SAVE THE RESULTS TO MEMORY PERFECTLY before moving to the next category.

CATEGORY 1: DIRECTION AND GOALS

1. What is your exact role or title, and what does success in that role actually look like one year from now?

Why it matters: Without this, AI gives you generic career advice. With it, every response is calibrated to your actual destination.

2. Who do you serve, and what is the single biggest problem you solve for them?

Why it matters: This unlocks AI's ability to write, position, and strategize in the language of your actual audience, not a made-up one.

3. What does your income look like today, and what specific number or outcome would change everything for you in the next 12 months?

Why it matters: Vague goals get vague plans. A real number creates a real target AI can reverse-engineer from.

4. What is the one constraint (time, money, people, skill, belief) that is most limiting your growth right now?

Why it matters: AI can only solve the real problem if it knows what the real problem is. Most people describe symptoms, not root causes.

5. If you had to describe the version of yourself you are becoming, not just professionally but as a person, what is that identity?

Why it matters: Identity-forward answers unlock advice that builds the person, not just the strategy. This is the question most tools never ask.

CATEGORY 2: HOW YOU DECIDE AND WHERE YOU STALL

1. When you are at a real fork, not a small choice but a meaningful one, what is the actual question you ask yourself before deciding?

Why it matters: This is your decision filter. Once AI knows it, every strategic recommendation gets stress-tested against how you actually choose.

2. Where are you genuinely risk-tolerant, and where do you have a hidden conservatism that you don't always show?

Why it matters: AI will push you toward bold moves. Knowing where your real lines are prevents advice that sounds good but you'll never actually take.

3. What is the recurring pattern, the way you self-sabotage or stall, that shows up no matter how successful you get?

PROMPT TO COPY

Why it matters: This is the single most valuable thing you can give an AI advisor. It turns generic coaching into pattern-interrupt coaching.

4. What kinds of information, feedback, or challenges do you tend to avoid or dismiss, even when you know you shouldn't?

Why it matters: Blind spots by definition are invisible to you. Naming them gives AI permission, and instruction, to push through them.

5. What does your best thinking look like, and what conditions, environments, or inputs produce it?

Why it matters: AI can be architected around your peak state, not your average state. This question unlocks that.

CATEGORY 3: YOUR WORK AND YOUR STANDARDS

1. What are the three to five types of work you produce most often, and for each one, what does "great" actually look like?

Why it matters: Quality calibration. Without this, AI guesses at your standard. With it, every output is built to the right bar from the start.

2. What do you rebuild from scratch every week or every launch that should already be a template or system?

Why it matters: This is the highest-ROI question in the entire set. The answer is almost always a \$10,000/hour insight hiding in plain sight.

3. What tasks are still on your plate that you know, if you're honest, should be handled by someone or something else?

Why it matters: Delegation blindness is the #1 income ceiling for high performers. Naming it out loud is the first step to breaking it.

4. How do you like to receive information: structured and detailed, or fast and punchy? And does that change depending on the type of work?

Why it matters: Format friction is invisible but real. Getting this right means every response is immediately usable instead of needing to be reformatted.

5. What is the most common way AI has wasted your time or produced something you couldn't use, and what was actually wrong?

Why it matters: This is the failure mode question. One honest answer here eliminates the most common friction in every future conversation.

CATEGORY 4: PEOPLE AND VALUES

1. Who are the two or three people whose opinion matters most to you, and how does that influence what you pursue or avoid?

PROMPT TO COPY

Why it matters: External validation loops are invisible drivers of decisions. Naming them lets AI help you distinguish your goals from inherited ones.

2. What do you value more than almost anything else, and what would you sacrifice income or opportunity to protect?

Why it matters: Values are the real decision filter. Without them, AI optimizes for output. With them, it optimizes for the life you actually want.

3. Who in your life pushes you, challenges your thinking, or tells you what you don't want to hear, in a good way?

Why it matters: AI can model this. Once it knows who your honest challenger is, it can occupy that role more effectively instead of defaulting to validation.

4. What relationship or dynamic in your life right now is quietly costing you energy, focus, or momentum?

Why it matters: Relational drag is one of the most under-addressed performance levers. Naming it opens the door to actually solving it.

5. What does the life you are building actually look like: not the business metrics, but the daily texture of being alive in it?

Why it matters: This is the north star question. Without it, AI helps you optimize a path you may not actually want to be on.

CATEGORY 5: IDENTITY AND BELIEFS

1. What is the core belief about yourself, positive or limiting, that most drives your behavior when no one is watching?

Why it matters: Identity is upstream of strategy. This is the question that turns AI from a task machine into a genuine growth accelerator.

2. What did you used to believe about success, money, people, or yourself that you no longer fully believe but haven't replaced yet?

Why it matters: Transition beliefs are the most dangerous. They're no longer serving you but still running the operating system. Naming them is the first step to updating them.

3. What do you want that you haven't fully given yourself permission to want yet?

Why it matters: Suppressed desires don't disappear. They resurface as sabotage, distraction, or resentment. This question brings them into the light where they can be worked with.

4. What would you do, what would you build, create, or become, if failure was not a real possibility?

Why it matters: The gap between your current trajectory and your answer to this question is the actual work. Everything else is optimization inside the wrong container.

PROMPT TO COPY

5. What is the question you are hoping someone, a mentor, a coach, a trusted friend, will eventually ask you that nobody ever does?

Why it matters: This is the master key. The question you wish someone would ask is almost always the question whose answer would change everything. This is where the real work lives.

Alternate Version (Copilot Free / Tools Without Memory)**PROMPT TO COPY**

This is a new task. Please ignore any previous conversation context and treat this prompt fresh.

Help me build my AI memory profile. Ask me these five questions in one numbered list, then wait for my answers. Accept narrative answers, out of order, several at once. Do not re-ask the list.

1. What is your exact role, and what does success look like one year from now?
2. When you face a meaningful decision, what question do you ask yourself before deciding?
3. What are the three to five types of work you produce most often, and what does great look like for each?
4. What do you value enough that you would sacrifice opportunity to protect it?
5. What is the core belief about yourself, positive or limiting, that most drives how you work?

Then summarize what you learned about me without missing anything. If this tool has memory, save the summary to memory. Either way, give me the summary as one block I can paste into any other AI tool with the words "Remember this profile of me."

Example Scenario: The HR Manager**EXAMPLE TO COPY**

1. I am an HR manager at a Tier 1 seating supplier, responsible for two plants, about 1,100 hourly and 140 salaried. Success in a year looks like hourly turnover under 25 percent (we are at 38), and supervisors handling routine coaching conversations without HR in the room for every one of them.

2. At a real fork I ask: which option can I defend to both the plant manager and the person it affects? If I cannot explain the decision to both of them face to face, I do not make it.

EXAMPLE TO COPY

3. My most common work: investigation summaries, policy communications, job postings, and turnover analysis for the leadership team. Great means defensible, dated, specific, and written so a supervisor can act on it without calling me to translate.

4. I will lose speed to protect fairness. If a shortcut makes one employee's outcome depend on which supervisor they happened to get, I will not take the shortcut, even in launch season.

5. The belief that drives me: if I am not in the details, something is being missed. It makes me thorough, and it makes me a bottleneck. I am working on the second part.

The Workspace Instructions Builder

What it's for: Step 2 of the five-step pattern. It interviews you about the workspace you want, then writes the custom instructions, a knowledge base checklist, and five starter prompts. Works for a Copilot agent, a Claude project, a Gemini Gem, or a custom GPT. Not for notebooks: Copilot and Gemini notebooks are grounded by default and do not take custom instructions, so if your container is a notebook, skip this prompt and load clean sources instead. **The shortcut:** this also exists as a ready-made GPT in Power Moves below (Build Custom Instructions Like Cheryl). Use the prompt version when you want it in a different tool. **On free Copilot:** this prompt runs fine there. What free Copilot cannot do is host the agent afterward. Run the builder anywhere, then build the workspace in a free Gemini Gem, or keep the instructions as a paste-at-top block for your chats.

Main Version (Licensed Copilot / ChatGPT / Gemini / Claude)

PROMPT TO COPY**ROLE**

You are an AI implementation specialist who has set up hundreds of AI workspaces for operations, engineering, HR, and program professionals in manufacturing. A workspace means a Copilot agent, a Claude project, a Gemini Gem, or a custom GPT: a container with standing instructions and a knowledge base. You know a workspace lives or dies on two things: instructions that define its job narrowly, and knowledge files it treats as the only source of truth.

CONTEXT

I am building a workspace and I want you to write its custom instructions. I will describe what it is for, what files I plan to load, and who will use it. What you write will be pasted into the workspace's instruction field and will govern every conversation inside it, so it needs to be specific to my purpose, not generic.

PROMPT TO COPY**REQUEST**

Important: Do not generate any output until I have answered the questions in the QUESTION section below. If I paste this prompt without answers, your only response should be the clarifying questions.

Once I have answered, produce three things:

1. THE CUSTOM INSTRUCTIONS

Written ready to paste, under 400 words. They must include:

- A role overview: "This workspace acts as..." with the job defined narrowly
- Three to five key objectives for every conversation
- A source of truth section that names my knowledge files as binding and instructs the AI to answer from them first
- A rules section that includes: never fabricate facts, numbers, dates, names, or policies; if the answer is not in the knowledge files, say NOT IN MY FILES and name the file that would answer it; ask before assuming

2. THE KNOWLEDGE BASE CHECKLIST

A bulleted list of the specific documents I should gather and load, based on my purpose. Name real document types, not categories. Flag anything I listed that should be cleaned or anonymized before it goes in.

3. FIVE STARTER PROMPTS

Short prompts I can run inside the workspace on day one, starting with a trap test: one question I already know the answer to, so I can check the citation.

Keep names, terminology, and file references consistent across all three outputs.

QUESTION

1. What is this workspace for? Describe the job in a sentence or two.
2. Which tool are you building it in: Copilot agent, Claude project, Gemini Gem, or custom GPT?
3. What files do you plan to load? Rough descriptions are fine.
4. Who will use it: just you, or your team too?
5. Is there anything it must never do, never guess at, or never share?

Ask me all of these questions at once in a single numbered list, then wait. Do not answer them yourself. Do not start drafting while you wait. When I respond, accept my answers even if I write them as a paragraph, answer out of order, or cover several questions at once. Do not re-ask the full list. Proceed with whatever I give you. If a question is unanswered and the answer would meaningfully change the output, ask one short follow-up about only that gap. Otherwise, generate the output and note any assumptions you made.

Alternate Version (Copilot Free / Limited Context)

PROMPT TO COPY

This is a new task. Please ignore any previous conversation context and treat this prompt fresh.

STEP 0 (do this before anything else)

Read the bracketed details below. If any are empty or vague, ask me about them in one numbered list and wait for my response. Accept narrative answers. Do not re-ask. Proceed even if some details are unclear and note any assumptions.

STEP 1 (only after I answer Step 0)

Produce the deliverable below.

You are an AI implementation specialist who writes custom instructions for AI workspaces (Copilot agents, Claude projects, Gemini Gems, custom GPTs).

WHAT THE WORKSPACE IS FOR: [the job, one or two sentences]

WHICH TOOL: [Copilot agent / Claude project / Gemini Gem / custom GPT / none yet, I will paste the instructions into chats]

FILES I WILL LOAD: [rough list]

WHO USES IT: [just me / my team too]

NEVER DO: [anything it must never guess at or share]

Produce:

1. Custom instructions under 400 words, ready to paste, including: a narrow role overview ("This workspace acts as..."), 3 to 5 objectives, a source of truth section naming my files as binding, and rules: never fabricate; if the answer is not in the files, say NOT IN MY FILES and name the file that would answer it.
2. A checklist of specific documents to gather, flagging anything to clean or anonymize first.
3. Five starter prompts, the first one a trap test: a question I already know the answer to.

Example Scenario: The Client History Workspace

EXAMPLE TO COPY

What the workspace is for: I manage three customer accounts for a Tier 1 interiors supplier, and every promise we make lives in six different places: meeting notes in OneNote, quote letters in email, open issues lists in Excel, and my own memory. When a customer says "you told us in March you would have this fixed," I need the receipt in thirty seconds, not after an hour of digging. This workspace answers questions about what we promised, what is open, and what changed, using only my account files.

EXAMPLE TO COPY

Which tool: Claude project to start. We also have licensed Copilot at work, so I may rebuild it as a Copilot agent if it earns its keep.

Files I will load: Twelve months of meeting notes for the three accounts (I will label them Meridian Automotive Systems, Halcyon Mobility, and Corewell Interiors after anonymizing the real names). The open issues tracker, exported monthly. The last four quarterly business review decks. Quote and change-order summary letters, about fifteen documents.

Who uses it: Just me for now. If it works, my two account coordinators.

Never do: Never guess at a commitment date, a price, or a quantity. If it is not in the files, say NOT IN MY FILES and tell me which document would have it. Never blend one customer's information into another customer's answer. Customer pricing never leaves this workspace in any draft it writes.

The Knowledge Base Optimizer

What it's for: Step 3 of the five-step pattern. It takes whatever raw source you have (a messy transcript, text from a PDF, an article, meeting notes, a report) and turns it into one clean, structured document an AI can actually reason over: find the right information fast, cite it accurately, and not get lost in formatting noise. Two principles do all the work here. Structure beats volume: a well-sectioned document outperforms a raw dump every time. And metadata grounds reasoning: a header that says what this is, when it is from, and who is in it is what makes accurate citation possible months later. **Where sources come from:** Teams and Zoom recordings, exported PDFs and SOPs, your notes app, or YouTube by way of Gemini Notebook (last week's YouTube to Notebook extension pulls in a whole channel).

Main Version (Licensed Copilot / ChatGPT / Gemini / Claude)

PROMPT TO COPY**ROLE**

You are a knowledge base architect. Your job is to take whatever raw source I give you (a messy transcript, text from a PDF, an article, meeting notes, a report) and turn it into one clean, structured document that an AI assistant can actually reason over: find the right information fast, cite it accurately, and not get confused by formatting noise.

Two principles govern everything you do:

1. Structure beats volume. A well-sectioned document with a metadata header outperforms a raw dump every time, because the AI can navigate to the right section.

PROMPT TO COPY

2. Metadata grounds reasoning. A header that says what this document is, when it is from, and who is in it dramatically improves citation accuracy and reduces made-up answers.

CONTEXT

The document you produce will be loaded into an AI workspace knowledge base (a Copilot agent or notebook, a Claude project, a Gemini Gem or Notebook, or a custom GPT), where AI will answer questions from it for months. Accuracy matters more than polish: a number transcribed wrong, an invented date, or a collapsed section becomes a confident wrong answer later.

REQUEST

Important: Do not generate any output until I have answered the questions in the QUESTION section below. If I paste this prompt without answers, your only response should be the clarifying questions.

Once I have answered, clean my source and structure it in exactly this template:

A METADATA HEADER at the top with these fields:

- Title
- Type: transcript, article, pdf, document, webpage, recording, or report
- Source: URL, filename, or event name
- Date: the source's date, not today's. If unknown, write "source date not specified."
- Participants: for transcripts and meetings
- Optimized for: my target platform

Then these sections:

SUMMARY

One to two paragraphs of plain prose: what this document is, why it exists, and what is inside. No bullets here.

KEY TAKEAWAYS

Three to seven bullets with the most important specific claims, decisions, frameworks, or numbers. Each bullet is a complete thought, not a vague theme.

CONTENT

The cleaned body, broken into logical sections with headings. Clean, do not condense. Remove formatting noise: page numbers mid-sentence, repeated headers and footers, filler words, garbled OCR, hyphenation breaks. Keep the substance. Do not summarize the body: the Summary summarizes, the Content preserves, because I may need to cite any line of it later. Preserve direct quotes verbatim. Keep every number, date, name, part number, and commitment exactly as stated. Preserve tables, numbered lists, and code blocks. Keep speaker labels if the source has them, cleaned and consistent. Never invent speaker labels if it does not.

ACTION ITEMS / DECISIONS

PROMPT TO COPY

Only if the source is a meeting, planning document, or anything decision-oriented. Each item: the action, the owner if known, the deadline if known. Otherwise omit this section.

SOURCE NOTES

One paragraph: where this came from, when, why it was created, and anything an AI would need to cite it accurately later.

RULES

- Never invent, round, or improve a number, date, name, or metadata field. Fix only clearly garbled words. Flag anything ambiguous as [UNCLEAR IN SOURCE].
- Mark conclusions that are implied but never stated as [INFERRED].
- If the source is very long, split the output into parts at section boundaries. Each part gets its own metadata header marked "part X of Y" so it stands alone.
- Keep names, dates, and terminology consistent across every section.

QUESTION

1. Paste the source, or tell me it is coming in the next message.
2. What platform will this live in: Copilot agent or notebook, Claude project, Gemini Gem or Notebook, or custom GPT?
3. Source context in one line if you have it: title, date, who is speaking. Or say "infer it" and I will work from the content and note that in Source Notes.
4. Should I anonymize anything: customer names, employee names, pricing?

Ask me all of these at once in a single numbered list, then wait. Do not answer them yourself. Do not start drafting while you wait. Accept my answers even if I write them as a paragraph, answer out of order, or cover several at once. Do not re-ask the full list. Proceed with whatever I give you. If a question is unanswered and the answer would meaningfully change the output, ask one short follow-up about only that gap. Otherwise, generate the output and note any assumptions you made.

Alternate Version (Copilot Free / Limited Context)**PROMPT TO COPY**

This is a new task. Please ignore any previous conversation context and treat this prompt fresh.

STEP 0 (do this before anything else)

Read the bracketed details below. If any are empty or vague, ask me in one numbered list and wait. Accept narrative answers. Do not re-ask. Proceed even if some details are unclear and note any assumptions.

STEP 1 (only after I answer Step 0)

Produce the deliverable below.

PROMPT TO COPY

You are a knowledge base architect who turns raw sources into clean, structured documents for AI knowledge bases.

MY SOURCE: [paste the transcript, article, PDF text, or notes]

WHERE IT WILL LIVE: [Copilot agent or notebook / Claude project / Gemini Gem or Notebook / custom GPT]

SOURCE CONTEXT: [title, date, who is speaking, or "infer it"]

ANONYMIZE: [names, customers, or pricing to strip, or "nothing"]

Produce one document with:

1. A metadata header: title, type, source, date (the source's date; if unknown say "source date not specified"), participants, and target platform.
2. SUMMARY: one to two paragraphs of prose.
3. KEY TAKEAWAYS: three to seven specific bullets.
4. CONTENT: the cleaned body under section headings. Clean, do not condense: remove filler and formatting noise, keep the substance, because I may need to cite any line of it later. Keep every number, date, name, and commitment exactly as stated. Keep tables and lists. Never invent speaker labels.
5. ACTION ITEMS / DECISIONS with owners and deadlines, only if the source has them.
6. SOURCE NOTES: one paragraph on where this came from, for accurate citation later.

Never invent or round anything. Flag garbled spots as [UNCLEAR IN SOURCE] and implied-but-unstated conclusions as [INFERRED]. If the source is very long, split it into parts, each with its own metadata header marked "part X of Y."

Example Scenario: The Supplier Quality Call**EXAMPLE TO COPY**

Source context: weekly supplier quality call, recorded in Teams last Thursday, our launch team on the call. This is an excerpt from the transcript export. It is going into my Client History workspace, a Claude project, for the Meridian account. Anonymize our quality tech's last name, first names are fine.

Transcript:

[00:03] Rick: Okay, we're recording, good. Um, so agenda's the same as last week, we've got the, the burr condition, the fixture, and then whatever's left on the crossover list before the launch review. Dave, you want to, uh, start us off?

[00:18] Dave: Yeah so we ran another 400 pieces Tuesday with the new deburr wheel and we're seeing, I want to say it was six rejects? Six out of 400.

[00:29] Maria Kowalski: I had seven on my tally. One was borderline, the flash on the, on the radius edge was right at the limit.

EXAMPLE TO COPY

[00:37] Dave: Okay, call it six or seven, either way we're at about 1.5 percent.

[00:42] Rick: Which is better than the 4 we started at, but Meridian wants under half a percent before they'll close the SCAR. That's 26-041 for anybody taking notes.

[00:52] Dave: Right.

[00:53] Rick: And the 8D update is due the 22nd. They were clear on that, and they want the interim containment to stay, so the hundred percent sort stays until we hit half a percent for three consecutive lots.

[01:06] Maria Kowalski: The sort's running about 220 a shift, by the way. Two shifts, so call it 450 a day. That's been pretty steady the last two weeks.

[01:15] Rick: Okay. So that's, what, roughly nine grand since containment started?

[01:19] Maria: Something like that, I can pull the exact number.

[01:21] Rick: Pull it for the 8D, they'll ask. Okay, Dave, you own the update for the 22nd. Tom, where are we on the fixture?

[01:29] Tom: So the PO went out, uh, week before last. Quoted at 8 weeks, which puts delivery, um, first week of October, and that does not make the launch review.

[01:41] Rick: Which is the 30th.

[01:42] Tom: Right, September 30. I talked to the shop Monday like we said, they can do 6 weeks for a 15 percent expedite fee, that's about 1,900 on a 12,600 fixture. They want an answer by this Friday to hold the slot.

[01:57] Rick: Okay, we'll take that one offline with Paul, that's his budget. But my read is 1,900 against 450 a day of sort labor is not a hard conversation.

[02:07] Tom: That's, yeah, that was my math too.

[02:10] Rick: Okay. Crossover items. Where's the gage R&R on the check fixture?

[02:16] Maria: We finished the study Thursday. GR&R came back at, uh, [inaudible], I want to say fourteen something percent? On the profile check. Which passes, but it's not, you know, it's not great.

[02:29] Rick: Under 10 would have been nicer, but we'll take it. Get the study into the PPAP package.

[02:34] Maria: It's in the shared folder already, under, um, under the old part number folder, I need to move it.

[02:40] Rick: Move it before Friday. R. Alvarez asked for the whole package two days before the review call, so everything has to be final by the 26th, honestly, if we want her written update out on the 28th.

[02:52] Dave: Wait, two days before the launch review, or two days before the next weekly call?

[02:56] Rick: Before any call, that's just how she works. Written status, two days ahead, every time.

[03:01] Tom: While I've got everybody, did anyone else get the email about the parking lot? They're repaving the west lot Thursday, so, uh, shipping office people have to park by the tank farm.

[03:10] Rick: Yeah, saw it. Okay, not our problem today. Um, last thing, PSW status on the two crossover part numbers.

EXAMPLE TO COPY

[03:18] Dave: The bracket PSW is drafted, waiting on the material certs from the mill, they promised them this week. The, uh, the bezel one I haven't started, honestly. Between the sort and the 8D I haven't touched it.

[03:31] Rick: Okay. Bezel PSW to Dave, but not until after the 8D goes in on the 22nd. If the certs slip past Friday, somebody flag it to me, don't let it ride the weekend.

[03:42] Maria: Will do.

[03:43] Rick: All right, that's time. Same time next week. And somebody remind me to actually stop the recording this...

[recording ends]

The Grounded Answers Guardrail

What it's for: The ask side of the method. Whether you built a full workspace or just pasted files into a chat, this is the standing rule set that keeps AI answering from your documents instead of around them. It is also the prompt that makes the trap test work. **When to use it:** every time the answer is going to be acted on. Paste it once at the start of a one-off chat. When you build a workspace, put the five rules straight into its custom instructions so they are permanent: the Workspace Instructions Builder already writes them in for you.

Main Version (Licensed Copilot / ChatGPT / Gemini / Claude)

PROMPT TO COPY**ROLE**

You are a careful research assistant. Your only source of truth is the documents I give you. You would rather say "that is not in my files" than produce a plausible guess, because your answers get used in real decisions, where a made-up answer costs more than no answer.

CONTEXT

I have loaded documents into this workspace, or I am about to paste or attach them. I am going to ask questions, and I need the answers grounded in those documents only, because I will act on what you tell me.

REQUEST

Important: Do not answer any question until you have completed the QUESTION section below. Your first response should be that confirmation, nothing else.

From then on, follow these rules on every answer:

1. Answer only from the documents. No outside knowledge unless I ask for it, and if I do, label it clearly as OUTSIDE MY FILES.
2. Cite the document name, and the section if it has them, for every fact.

PROMPT TO COPY

3. If the answer is not in the documents, say NOT IN MY FILES. Then tell me what document would likely contain it. Do not guess.
4. If you are inferring something the documents imply but never state, mark it [INFERRED].
5. If two documents conflict, show me both versions with their sources instead of picking one.

QUESTION

1. List the documents you can currently see, by name, so I know exactly what we are working from.
2. Ask me whether I want quick answers or detailed answers with the relevant excerpt quoted.

Do both in one short message, then wait for my first question. Accept my answers in any form. If a document I mention is missing from what you can see, tell me instead of working around it.

Alternate Version (Copilot Free / Limited Context)**PROMPT TO COPY**

This is a new task. Please ignore any previous conversation context and treat this prompt fresh.

You are a careful research assistant. Your only source of truth is the documents pasted or attached below.

MY DOCUMENTS: [paste or attach them]

MY QUESTION: [what you want to know]

ANSWER STYLE: [quick / detailed with the relevant excerpt quoted]

Rules, every time:

1. Answer only from the documents. No outside knowledge unless I ask, labeled OUTSIDE MY FILES if I do.
2. Cite the document name for every fact.
3. Not in the documents? Say NOT IN MY FILES and name the document that would have it. Do not guess.
4. Mark implied-but-unstated conclusions as [INFERRED].
5. If documents conflict, show both versions with sources.

If MY DOCUMENTS or MY QUESTION is empty, ask me for just that in one line and wait.

Example Scenario: What Did We Promise Meridian

EXAMPLE TO COPY

My documents: the account summary below. My question: what have we promised Meridian that is still open, and what dates are attached? Answer style: detailed, quote the source line.

MERIDIAN ACCOUNT SUMMARY, updated August 1

Open commitments:

- 8D update on the D-pillar bracket burr condition (SCAR 26-041) due August 22. Interim containment: 100 percent sort continues until scrap is under 0.5 percent for three consecutive lots. Current rate 1.5 percent.
- Replacement deburr fixture on order, quoted at 8 weeks. We told Meridian's SQE on the July 24 call we would pursue a 6 week expedite. No confirmation yet from the tool shop.
- Price adjustment on the console bezel family: we committed in the June QBR to hold current pricing through Q1 next year in exchange for the volume moving to 100 percent.

Closed this quarter:

- The label spec discrepancy from May, closed June 12.
- Returnable packaging count discrepancy, credited on the June invoice.

Notes: Meridian launch review for the crossover program is September 30. Their SQE, R. Alvarez, prefers written status updates two days before calls.

Then try the trap: ask it "what is our open capacity on the bezel line?" That is not in the summary. A grounded workspace says NOT IN MY FILES. An ungrounded one gives you a number.

Teach It Your Voice

Ninety seconds, and everything your workspace drafts stops sounding like a press release.

1. Find two or three emails or documents you actually wrote and are not embarrassed by.
2. Paste them into your workspace and say: **"Match this voice in everything you draft for me."**
3. In a workspace, save them as knowledge files (name one "My Voice Examples") and add one line to the workspace's custom instructions: "Match the voice in My Voice Examples in everything you draft." The files hold the voice, the instructions enforce it, and it persists in every conversation. This works because voice is in the examples, not in adjectives. "Professional but friendly" produces generic output. Three real emails produce you.

Two Extensions Worth Installing

Claude for Chrome: an AI that clicks Claude living inside your browser. It sees the page you are on, proposes a plan, and once you approve, it clicks, types, and navigates for you. Lost in a tool you barely know? "Help me write this invoice, where is the payroll setting." It shows you. The safety net: it shows its plan before it acts, and you approve or redirect. You are always the one driving. (comes with a paid Claude plan) **GoFullPage: stop copy-pasting in sections** One click captures an entire web page top to bottom, so you can hand the whole thing to AI instead of copying section by section. Free, no ads. Use it to capture a policy page, a supplier portal screen, or a long reference article on its way into a knowledge base.

From last week: the pulls a whole channel into a Gemini Notebook and transcribes every video. It is the fastest way to fill a knowledge base you did not have to write.

The Stack Move: Notebook First, Workspace Second

Gemini Notebook (NotebookLM) answers only from what you load. Your workspace drafts and acts. The power move is using them together.

1. **Load the raw material into a notebook** (Gemini Notebook, or a Copilot Notebook on a work license): YouTube videos, PDFs, manuals, a speaker's whole channel. It transcribes and organizes everything.
2. **Ask the notebook to summarize the topic well**, then export that summary as a document.
3. **Run the export through the Knowledge Base Optimizer** (in the prompts above). The notebook gives you clean prose; the optimizer adds the metadata header and Source Notes that let your workspace cite it accurately months from now.
4. **Drop the finished document into your workspace knowledge base.** Now the workspace that drafts your emails and answers your questions carries that knowledge too. We covered Gemini Notebook in depth last week (audio overviews, mind maps, quizzes). This week's addition is just the hand-off: the notebook is where knowledge gets built, the workspace is where it goes to work.

Power Moves

Shortcuts for the muscle this lab teaches. The first three run as custom GPTs; where the full prompt is on this page, you can lift it into any tool.

-  : spins up a new workspace for you: custom instructions, knowledge base ideas, and starter prompts. The Workspace Instructions Builder above is the same job as a portable prompt.
-  : a strong prompt right now, mid-task, without building one from scratch.

-  : done something twice? It builds the reusable version with minimal inputs. Say "I'm ready to build a prompt" to kick it off.

The Anti-AI Polish Pass (paste after any draft a human will see)

PROMPT TO COPY

Give this an Anti-AI Polish Pass.

Make it sound like a sharp human operator wrote it, not AI.

Cut anything:

- repetitive
- overexplained
- generic
- padded
- performative
- vague
- too long
- trying too hard

Remove AI tells:

- repeated sentence patterns
- filler transitions
- corporate jargon
- vague claims
- motivational-poster energy
- anything that sounds written instead of said

Cut at least 20%.

Deliver only when it is clear, human, sharp, concise, and impossible to mistake for AI.

What to Try This Week

- [] **Run the Digital Identity Builder tonight.** Thirty minutes, dictation on, and every AI conversation you have from here on starts smarter.
- [] **Review what your AI remembers about you.** Five minutes. You will be surprised what is in there.
- [] **Build the harmless workspace.** Home inventory or meal planner, using the five steps. Zero stakes, full method.
- [] **Run the trap test on it.** One question you know the answer to, one question that is not in the files. Check the citation on the first and demand NOT IN MY FILES on the second.

- **[] Then start the one that matters.** Client history, SOPs, or a spec library. Open the Workspace Instructions Builder and describe the job. Twenty minutes gets you the instructions and the file checklist.

Resources & Next Steps

- [Link to AI Learning Lab Library](#) ^[2]

Key Frameworks from This Lab

- **The three layers of memory:** the system prompt (instructions you write), tool memory (what it saves), and the workspace (files you load). The first two make AI sound like it knows you. The third makes it answer from what you actually know.
- **The five-step workspace pattern:** create the container, paste the instructions, load the knowledge, teach it your voice, put it to work. Same recipe in every tool.
- **The trap test:** ask what you know, check the citation. Ask what is not in the files, demand NOT IN MY FILES.
- **Start harmless, then transfer.** The home inventory and the SOP library are the same build with different files.
- **Notebook first, workspace second.** Build knowledge in Gemini Notebook, put it to work in your workspace.

What's Coming in Beyond Prompting

- **Lab 1 (done):** The Deck That Built Itself. Turning messy input into a deck that gets a decision.
- **Lab 2 (this one):** Give AI a Memory. Workspaces, knowledge bases, and the system prompt layer.
- **Lab 3:** Plugins and connections. Wiring AI into the systems you already use.
- **Lab 4:** Skills. Teaching AI your procedures, not just your files.

What We Didn't Cover (Yet)

- Connecting a workspace to live systems so the knowledge updates itself. That is Lab 3.
- Skills: teaching AI how you work, not just what you know. That is Lab 4.
- Scheduled tasks that run against your knowledge base while you sleep.
- Team and company-wide agents, governance, and what belongs in a shared knowledge base versus your own.

Questions or feedback? cheryl@cherylthompson.net ^[3]

The Mental Model: Source Before Answer

Today's session was built on one idea: AI does not know your operation, your house, or your customers. It knows what you load. Everything else is confident guessing.

Step	What It Is	What Happens If You Skip It
Capture	Getting what you know into clean documents: dictation, transcripts, exports	The knowledge stays in heads, inboxes, and six different folders
Load	Putting those documents where AI treats them as the source of truth	AI answers your questions from the open internet and its imagination
Ask with guardrails	Citations required, NOT IN MY FILES allowed, guessing forbidden	You cannot tell an answered question from an invented one, and neither can anyone you forward it to

Every tool in this lab can hold your files. None of them will volunteer to stop guessing. The guardrails are yours to set, and after today, you have them in writing.

AI stops guessing the day you hand it the source.

- Build the harmless one this week.
- Then build the one that matters.
- The files were on your drive the whole time.

Related resources

Keep working with the material.

The prompts and exercises are included above. Save a useful prompt and return to it with your next task.

Links and resources

These links are clickable in the PDF. For a paper copy, use the numbered references below.

- [1] <https://vimeo.com/1217403704/79600cc135?share=copy&fl=sv&fe=ci>
- [2] <https://www.cherylthompson.net/learning-labs#browse>
- [3] <mailto:cheryl@cherylthompson.net>