

YOUR JOURNEY — what will happen, in order

This guide is for you. The other file, **EARNED TRUST PROTOCOL**, is for your AI: you paste it, you never have to read it. To begin: open the AI you'll write with, paste that whole file into the chat, and say hello. It will greet you and offer three choices. If it does, it's ready. If it just says "OK," paste again — some AIs need the whole file in one go. (If your AI lets you attach files, you can attach the protocol instead of pasting it.) Your AI takes it from there. Then come back to this guide whenever you want to know what's coming.

What you're about to do, in one paragraph

You are going to make something — an essay, a cookbook, a family history, a research paper — with an AI helping, and you are going to keep a receipt of exactly how it helped. Not from memory, weeks later: as you go, at the end of every session, in a plain little logbook called your **00 INDEX**. Then you'll have your draft attacked by a second AI on purpose, check every source with your own eyes, say plainly on the finished work what the AI did, and put the work somewhere permanent, where anyone can find it and check it. That's the whole thing. You stay the author. The AI does the bookkeeping.

Two things to know before you start

Your writing AI has the whole method, and you don't have to remember any of it. If it ever seems to have lost its place, say **"where am I?"** — and if it still doesn't do the next thing, tell it the step by name: *"we're at Step 3 — please build the critic packet."*

Three people are in this story: **you**, **your writing AI** (the one you think and draft with), and **the other AI** (a critic from a different company, whose only job is to attack your drafts).

Setup, then the eight steps

Each step says who does what and what you'll do with your hands — type, copy, paste, open, click. Your writing AI announces each step as you arrive ("Step 2 of 8 — Draft").

Setup (about ten minutes, once). You tell your writing AI what you're making. At the first session your AI tells you which of three things it can do where you are: make files you download, open your own files, or only give you text. If text is all it can do, it hands you the words and you make the folder yourself next time you're at a computer — that's normal. Either way, by the end of setup you have a folder and a logbook, and your AI has taught you the sentence that ends every session (it's near the bottom of this guide, under "The sentence you say to end a session"). Already have writing — a thesis, a manuscript, years of notes? Tell your AI at the start. It fingerprints what you have on day one (a fingerprint — a short code that proves the file hasn't changed since), calls it your Legacy Work, and together you look at what's already written and decide what drafting is still needed. The record still starts today; it just starts with a photograph of what already exists.

Step 1 — Choose your team

Two decisions, and your writing AI walks you through both in plain words. The first is your **writing partner**: the AI you'll think and draft with, session after session. The second is your **critic**, or critics: one or more *other* AIs whose only job is to attack your drafts at Step 3. The names are examples, not a shopping list — Claude, ChatGPT, Gemini, Perplexity, Grok, Copilot — or whichever you already have. More is fine; some people use four.

The critic should come from a different company than your writing partner, because two windows of the same AI share the same blind spots. Only one AI? That's all right: a brand-new conversation with the same AI will do the critic's job, and your logbook will say "same service" — honest, a little weaker, and fine. For work that makes checkable claims — a research piece, anything with sources — a second company is required, not optional; a free account there is enough. You say which is which; your writing AI writes both choices into your logbook.

Step 2 — Draft, with your partner

This is the long step, and it's the partnership step. You write. You ask your AI what it thinks. You take some of what it says and leave some. When it has your voice wrong, you tell it — "make it mine again" — and it does. Early on it will show you one paragraph three ways and ask which sounds like you; that's it learning your ear. You're getting to know each other: how it helps, where it drifts. Session after session. Every version is saved, numbered, never written over, so nothing you liked is ever lost.

One warning, said kindly. When your writing AI tells you everything is wonderful, that's what it does — it isn't evidence. Step 3 exists for that.

Step 3 — Have it attacked

In order:

1. **Your writing AI writes the whole message for you** — your draft, plus a request to be a tough critic — as one piece of text.
2. **You copy that text.** (Too long to paste in one go? Say so — your AI prepares a file you can attach, or breaks it into numbered sections; it never quietly leaves part of your draft out.)
3. **You open the other AI** — your critic from Step 1. (Only one AI? Open a brand-new conversation with it.)
4. **You paste the text in and send it.**
5. **The other AI writes back.**
6. **You copy its reply, come back to your writing AI, paste it in, and say "here's what it said."**
7. **Your writing AI saves the reply word for word** and helps you decide what to do about it — that's Step 5.

How often? Once per finished piece — and "a piece" is whatever you and your writing AI agreed at the start: a chapter, a recipe with its story, the whole thing. You say when

it's done. Over a whole project that's a few critiques, an afternoon each — not every week. If the work is long enough that the pieces must hang together, one more of the whole at the end; your writing AI will say whether that applies to yours.

You can also ask for one whenever you like — on a paragraph, an idea, a chapter. Say “I’d like a critique of this” and your writing AI builds the message exactly as above. Extras never replace the required ones, and they read well in the record, not badly. This step is never skipped.

Step 4 — Check every source yourself

A **source** is anything your work says is true because someone else said it — a book, an article, a website, a record. *What your writing AI does:* lists each claim and where it came from, one at a time — never a giant table. *What you do:* open the source yourself and read the part that matters, then tell your writing AI “yes, it says that” or “no, it doesn’t.” It writes your answer down beside the claim. Nothing is checked for you; the “I opened it and it says this” part is yours alone. One at a time, until the list is done. Never skipped.

“I can’t check this” is a fair answer — your AI writes the claim down as unverified, and you come back to it or leave it marked. Writing a memoir or a story? Then most of this step doesn’t apply — only the things you present as fact get checked, and your AI notes why the step was short.

Step 5 — Decide what to keep

The partnership again. The two of you go through what the critic said, point by point. For each one you decide: take it, or say why not — both answers are fine. Your writing AI writes down each decision and why, so a reader later can see what changed, what didn’t, and your reason.

Step 6 — Keep the receipts

The logbook row is one line in your **00 INDEX** — the logbook — saying what happened this session: the date, which AI, what got done. Your writing AI writes it when you say “**log it and close out.**” This step runs the whole time, at the end of every session. To see your logbook, ask — your writing AI shows it to you — or open the file called **00 INDEX** in your record folder.

Step 7 — Say how it was made

Two small things go on the finished work. **The note on the work** is a few honest sentences on the work itself: which AIs you used and what each did. **The record card** is a QR image — the square barcode a phone camera reads — that a reader scans to see your logbook. Your writing AI drafts both and shows you; you say yes, or change what you’d like changed. Four small layout choices come with it — where the mark (the small label saying AI helped) sits, a short explaining paragraph on the cover or not, whether the mark carries the LI- prefix (a code for work that existed

before the record began), and where the card is placed. Your writing AI suggests an answer for each, in one breath; one word from you (“fine”) accepts them all.

Then two questions your AI asks **on their own**, one at a time, because each one decides what becomes public. First, a contact address or none: your AI shows you exactly what would be printed, so you can see what a stranger would see. Second, the locked box, in these words: *“The public deposit is your work, your record card, and your fingerprint list. Would you also like to publish a locked, passphrase-protected copy of your working records?”* Each gets its own answer from you; neither is folded into the “fine.”

Step 8 — Publish

The computer day. Before you start the upload, open **5 WHAT ZENODO WILL ASK YOU.pdf** in the folder **1 READ THESE FIRST (PDF)** — it walks Zenodo’s boxes one by one. If you chose the locked box at Step 7, follow **4 LOCK YOUR RECORD** first and confirm the box opens for you before the check. Then, in order:

1. **Copy your whole record folder somewhere safe first** — a USB stick, your cloud drive, a second place on your computer. If you ever lose a passphrase, that copy is your record.
2. **Your writing AI seals the record** — one command that writes a fingerprint of every file going up — **and runs the check** on those files, then reads you the verdict word for word: **safe to upload**, or **stop**, and why.
3. **You open the folder called 05 Deposit** inside your record folder and upload exactly what’s there — nothing more, nothing less.
4. **You fill in Zenodo’s boxes** (Zenodo is the free archive your work goes into) by copying from the document called **3 PASTE THIS**, which your writing AI has already written for you.
5. **Before you press Publish, preview the page.** Zenodo shows you what the public will see. Check five things — your name, the title, the contact address (or that there is none), the licence, and the list of files — and read them back to your AI. If anything is wrong, fix it now; after Publish, nothing comes down.
6. **You press Publish.**
7. **After you press Publish, check it.** Open the public page. Open your work from it. Scan the record card with your phone (or open the record link) and confirm the record it opens is *your* record — this project, this logbook. Then copy the address you’re given — the **DOI**, a permanent address for your work, the kind libraries use — and save it in your logbook: your AI writes it into the Deposited row.

You’re finished when you’ve checked it, not when you’ve been given an address. Your work, your record card, and the seal (the fingerprint list) are every deposit. The **locked box** — a passphrase-protected copy of your drafts and transcripts — is the one optional thing.

Where you can do each part

Sessions — a phone is fine. Steps 1 through 6 are conversations: talk, paste, decide.

Step 7’s card and Step 8’s locking, checking, and uploading — sit at a computer. Those want a real keyboard and a real file folder. Plan one afternoon at a desk, near the end.

One honest note about phones: from a chat window on a phone, your AI usually can’t touch files — it’s the “only give you text” case from Setup. It hands you the text — the logbook row, the folder plan — and you make the folder and paste the text in next time you’re at a computer. That’s normal; the record is still yours and still complete.

How long things take

(These are the founder's estimates from his own projects — your project is yours.)

- **Setup:** about ten minutes, once.
- **A light project** (a short essay, a family story, a talk): 4 to 6 hours of your time, spread over 4 to 6 weeks.
- **A heavier one** (a paper, a book chapter, anything with many sources): 15 to 50 hours, over 8 to 12 weeks.
- **Publishing day:** about an hour, of which the Zenodo upload itself is about fifteen minutes. The rest is copying, checking, and reading a document your AI wrote for you weeks ago. The weeks matter more than the hours: sources take time to open, and critics take time to answer.

What finished looks like

On the day you publish, your work folder holds three numbered things, in the order you use them, and nothing else:

Name	What it is	What you do with it
1 READ FIRST/	What to do today, in order — your AI wrote it	Open it first and follow it
2 YOUR RECORD - stays home/	Your logbook and evidence	Keep it; this folder never leaves your computer
3 PASTE THIS/	The words Zenodo will ask for, ready to copy	Copy from it, box by box

Inside 2 YOUR RECORD - stays home, one folder called 05 Deposit holds the only files that go up: **your work**, **your record card** (the QR image), and **the seal** — MANIFEST - seal <date>.md, a fingerprint of each file, so anyone can later prove nothing changed. Those three are required; together they are a complete deposit. If you chose to publish your working record too, a fourth file sits beside them: the **locked box**, a .7z that opens only with a passphrase only you hold. The locked box is the one optional thing.

Then you press Publish, Zenodo gives you a **DOI** — a permanent address for your work, the kind libraries use — and you open the page, open the work, scan the card, and save the DOI in your logbook. That check is the finish line.

The sentence you say to end a session

“Log it and close out.”

Say it and the AI writes the logbook row, refreshes your record link (your project's save-code, so you can pick up anywhere), and hands everything back to you. Before

you leave, it makes sure today's draft and the updated logbook are saved where you can reach them: if it can make downloads, it gives them to you and asks you to open one; if it can only give text, you copy it into your own notes or document app now and tell it when it's saved. Say it every time, even after ten minutes — short sessions are the ones people forget to log. "I'm done for today" works just as well. What happens next time is the next section of this guide, "How to come back tomorrow."

How to come back tomorrow

Every time you close out, your AI gives you a record link — your project's save-code, a long web address that carries your whole logbook inside it. Keep it somewhere you'll find it (a note, an email to yourself). Next time, bring two things: the record link **and** your latest draft or your saved folder. The link is the logbook — it tells your AI what has happened. The files are the work — a fresh conversation has no memory of your manuscript, so without them it can only tell you where you were, not carry on. Open a new conversation, paste the protocol again if the AI has forgotten it, paste your record link, give it the draft, and say "let's continue." Your AI tells you the last finished step and asks you to confirm what comes next. If your AI can see your files, you can skip the link: say "continue my project" and it opens your shelf and asks which one.

What you'll need

An AI on a plan that can hold a long conversation (the protocol alone is about twelve thousand words, and the work runs for weeks); free plans often can't, a paid plan usually can; a month or two usually covers a project — check your plan's cancellation terms before you start; the time estimates above are examples, not a schedule. A second AI service for the hostile read — a free account there is fine. If you're doing this as part of the Earned Trust Beta Trial: taking part may cost a month of a subscription; the study pays nothing and asks for nothing.

Only one AI? A fresh conversation with the same AI can do the hostile read at Step 3. Your record will say "same service" — honest, a little weaker, and that's fine.