

The PII-Safe List Prep Kit

A reconstruction of AJ Wilcox's workflow for cleaning client contact lists without sending personal data to a cloud model · The AI Social Playbook

What this is, and what it isn't. AJ dictated his instructions live on camera rather than pasting a saved prompt, so the panels below reproduce what he said, tidied into copy-pasteable form and filled out where a sentence trailed off. The agent's confirmation reply was summarised on screen and is not reproduced. A few hardening lines were added and are marked as such. This also is not a security guarantee: it reduces exposure and creates a moment where a failure becomes visible, which is not the same as making the data incapable of leaving.

01 Technical Stack

TOOL	ROLE IN THE WORKFLOW	NOTES
OpenClaw	The agent harness. Receives your instructions, decides which model handles each step, and runs the file operations.	Open source and free. Needs a machine of your own to live on. AJ runs his on a Mac Mini in his basement.
Telegram	The interface you talk to the agent through. One channel per agent.	AJ's channel is named Luna; the name is arbitrary and you address the agent as OpenClaw. Note that files you send do travel through Telegram on the way to your machine.
Anthropic Claude	The orchestrating cloud model. Reads your instructions and drives the workflow, and by design never reads the client data itself.	AJ names Sonnet in the demo. Any cloud model the harness supports can hold this role.
Ollama	Runs the local model on your own hardware, so the file contents stay on the machine.	Installed on the same machine as the harness.
GLM 5.2	The local model that actually reads and maps the client data.	AJ's current choice. He states directly that any model you have available locally will do the job.

LinkedIn Campaign Manager	Supplies the contact-list and company-list templates, and receives the finished CSV as a matched audience.	Plan, then Audiences, then create a matched audience.
Speech-to-text dictation	Optional. AJ dictates his instructions to the agent rather than typing them.	Not required by the workflow. Any dictation tool works.

02 Configuration

Have all four in place before you send the first message.

A machine that stays on

The harness and the local model run on hardware you own, so something has to be awake when you send a file. AJ uses a Mac Mini kept in his basement. This is the real cost of the workflow and there is no cloud substitute for it.

Ollama installed with at least one model pulled

The cloud model can only hand the data off if there is somewhere to hand it to. Pull the local model before the first run rather than during it.

LinkedIn's templates, stored with the agent

Download both the contact-list and the company-list template from Campaign Manager, then give them to the agent once in a separate conversation with the instruction to store them and reuse them for every future list. AJ did exactly this, so he never has to re-explain the column structure.

A sanitized test list ADDED STEP

Run the whole sequence once on a list with fake names and addresses, and confirm the agent honours the handoff, before you point it at a real client's data. AJ sanitized the list he demonstrated with for exactly this reason.

03 Flowchart Map

1	Store LinkedIn's templates with the agent IN Contact-list and company-list CSVs from Campaign Manager → OUT Templates the agent reuses for every future list
2	Send the rule, with no data attached SYSTEM PROMPT 1 IN A description of the job plus the PII restriction → OUT The agent's written confirmation of the terms
3	Read the confirmation before you attach anything IN The agent's reply → OUT A go or no-go decision, made by you, while the data is still on your machine
4	Send the client workbook SYSTEM PROMPT 2

	IN The messy multi-tab Excel file → OUT A record count and sheet breakdown, reported by the local model
5	Send the template and the build instruction SYSTEM PROMPT 3 IN LinkedIn's contact-list template → OUT One mapped CSV per source sheet, emailed to you
6	Restructure the rows for match rate IN The mapped CSVs → OUT One block of email-only rows, then a block carrying first name, last name, title and company
7	Upload as a matched audience IN The restructured CSV → OUT A match rate returned in 48 to 72 hours, typically 60 to 90 percent
8	Convert the contact list into a company list IN The same cleaned data → OUT An ABM list for targeting named accounts, or for excluding customers you already closed

04 System Prompts

Send these as three separate messages, in order. Do not attach the data file until the agent has answered prompt 1. The final line of prompt 1 and the record-count request in prompt 3 are hardening added on top of what the demo showed.

SYSTEM PROMPT 1

The Handoff Rule

COPY WHOLE

Here is what I want us to do. I am going to give you an Excel sheet with **[NUMBER OF TABS]** different tabs on it. These are a contact list from a client. I am also going to give you the template file that LinkedIn provides, which is a CSV.

What I want you to do is take the data from the actual data file and apply it into the template, in the structure LinkedIn expects.

One caveat. Because we are dealing with PII, I do not want you as **[CLOUD MODEL NAME]** to touch any of this data. I want you to spin up an Ollama local model to actually deal with the data. Use **[LOCAL MODEL NAME]**, or another local model we have access to.

Make sure that you fully hand this file off. When I give it to you, do not ingest it at all. You only send it to the local model to read.

Let me know if you are okay with this, and restate the restriction back to me in your own words. Do not ask me for the file until you have done that.

SYSTEM PROMPT 2

[COPY WHOLE](#)

The Data Handoff

This is the client data. It is a contact list on **[NUMBER OF TABS]** separate tabs of this Excel file, so our output is going to be **[NUMBER OF TABS]** different files.

Make sure this goes directly into the local model. You as **[CLOUD MODEL NAME]** do not touch this.

Let me know when you are ready and I will give you the template file as well.

SYSTEM PROMPT 3

[COPY WHOLE](#)

Build and Deliver

Here is the template for LinkedIn contact lists. Store this somewhere you can access it. Every time I give you a list from now on, you are going to use one of these.

Go ahead and build out **[NUMBER OF TABS]** CSV files using the contact lists from the previous file. Map each source column into LinkedIn's template structure: email address, first name, last name, title, company name, country.

Before you build anything, tell me how many records you found and how they are distributed across the sheets, so I can check that against the source.

Leave a cell empty rather than guessing at a value. Do not invent a job title or a company name that is not in the source data.

When you are finished, email a copy of all **[NUMBER OF TABS]** files to **[YOUR EMAIL ADDRESS]**.

05 Variable Callouts

VARIABLE	WHAT IT IS	WORKED EXAMPLE
[NUMBER OF TABS]	How many sheets the client's workbook contains, which is also how many output CSVs you get	three
[CLOUD MODEL NAME]	The orchestrating model you are telling to keep its hands off the data, named explicitly so the instruction has a subject	Anthropic Sonnet
[LOCAL MODEL NAME]	The model you have pulled under Ollama on your own machine	GLM 5.2

[YOUR EMAIL ADDRESS]

Where the finished CSVs get delivered,
so you open them from your inbox
instead of hunting through the host
machine's downloads folder

you@youragency.com

06 Human Quality Control

Three checks before anything ships. Each one catches a failure observed in the episode.

Check 1. The agent restated the handoff restriction before you attached the file, and you read what it said rather than scanning for the word yes.

CATCHES The agent agreeing in general terms and then processing the data itself. This is the specific failure AJ described: it said it would spin up a local model, then came back later and admitted it forgot. By that point the cloud model had already read the file.

Check 2. The record count the agent reports matches the row count in your source workbook, sheet by sheet.

CATCHES A tab skipped, a header row counted as data, or rows dropped during the mapping. In the demo the agent reported 986 records across three sheets, which is the number you would check against the file before letting anything upload.

Check 3. The finished CSV carries first name, last name, title and company, and any email-only rows sit in their own block rather than mixed through the file.

CATCHES A match rate around 50 percent instead of 90 percent or better. People hand out a work email and log into LinkedIn with a personal one, so an email-only upload matches roughly half the list. The four identity fields are what let LinkedIn resolve a single person.
