AI POWER HOUR · SESSION 2 OF 4 · COMMUNITECH

Hiring your first 5 AI employees

What we covered, so you can build your own small team of agents at home.

November 4, 2026 · Heather Gawel and Babajide Okusanya, Momentum Product Co. ·
Sponsored by Google for Startups

1. One job each

A skill is a repeatable job. An agent is a team member with one job, clear rules and limited reach.

2. Files live outside

Keep your files outside the agent, with tight read and write permissions.

3. Talk, don't type

Don't write your own instructions. Talk it through, let the AI write them, and fix them every time they fail.

4. Draft, don't send

Never let an agent pay, send, reply or forward anything until you trust it.

1 A scheduled task isn't an agent

Both run on their own. Only one has limits.

A scheduled task is a prompt that runs at a set time. That covers a lot of useful automation, like daily follow-ups, a morning focus brief or a Friday wrap-up. An agent goes further. It has:

- **A specific job**, and only that job
- **Rules** for how to do it, including when to stop
- **Limited reach**, so it can only read, write or change the files and tools you've given it

Control is the real difference. With one big scheduled chat, it's hard to limit what it can touch.

2 Skills, connectors and plugins

Three building blocks, and most people only use the first one.

- **Skills** are instructions for anything you do over and over. If every invoice gets checked against its quote and logged to a sheet, that's a skill.
- **Connectors** link the AI to your stuff: Google Drive, calendar, email, your CRM.
- **Plugins** (in Claude Cowork) bundle several agents and skills into one package you upload and run. Most people haven't found them yet, and they're where the agent capability really lives.

Stacking skills in one chat isn't quite an agent. When everything sits in one place, everything breaks at the same time.

3 How we built the demo team

A bookkeeper, a legal reviewer and a chief of staff, each with its own file.

The **bookkeeper** checks invoices against quotes. The **legal reviewer** looks at incoming contracts. The **chief of staff** doesn't do the work. It decides who does it, and brings you only the things that need your decision.

- **One instruction file per agent.** Each agent gets its own .md file, plus a CLAUDE.md that explains who does what (your basic operating instructions).
- **Agents don't hold files.** Invoices, quotes and the expenses sheet live in a separate folder you control. The agent reads from it and writes to it.
- **Spell out read, write, change and share.** The bookkeeper can read invoices and quotes, but the expenses sheet is the only file it's allowed to change.
- **Spell out the job step by step.** Find the matching quote. If there's no quote number, match on vendor and description. Check the math, including HST. Compare totals. Skip anything already logged.
- **Add decision logic and limits.** Mark it approved or disputed, add a note, draft a short reply to the vendor for disputes, then report back. Never pay, send, reply or forward.
- **Run it once, then schedule it.** We zipped the files, uploaded them as a plugin, ran it, checked the output, then set it to run every day at 8 AM. Because it runs unattended, it writes its questions to the folder instead of waiting for an answer.

Three invoices, three decisions. One was over its quote, one looked like a duplicate we'd already paid, and one had a delivery fee that wasn't on the quote. You reply on your phone, in Slack or by email, and it moves to the next step.

4 Don't write your own instructions

Good skill and agent files are long. Nobody should be typing them from scratch.

- Talk it through out loud. Describe how the job is done in your business and ask the AI to turn it into a skill.
- Ask it to coach you: what does the output look like, what are the headings, what are the gaps?
- Then ask, "What did I miss?" before you paste anything in.
- Build the instructions in a regular chat, not in the agent's instruction box. The box works, but the results are usually weaker.
- It's fine to use one tool to write instructions for another, like Claude writing instructions for a ChatGPT agent. A tighter prompt also saves tokens.

It will fail, and that's the point. Every time it does, update the .md file. These files get better every time you fix them.

5 Draft, don't send

Client communication is worth a quick review.

If you're new to building agents, don't let them send email. Have them draft instead. Voice dictation picks up whatever's said in the room, and it can end up in a client email you didn't read first. Low-stakes replies, like confirming you got a support ticket, might be fine.

A dispute over an invoice isn't low stakes. Your comfort level is your call, but start with drafts.

6 Claude Cowork vs. ChatGPT

Both can do this. They're set up differently.

- **ChatGPT Plus** is great for GPTs and single skills, but it struggles to coordinate several agents. If you're on Plus and want orchestration, pair it with another tool.
- **ChatGPT Business** costs a little more and adds the full set of connectors (GitHub, Granola, Zoom and more), separate connectors for each agent, agent-owned accounts with their own drive and email, and much finer permissions. It's also a more comfortable place for customer data.
- **Claude Code** in the terminal gives you the most freedom, if you're comfortable there.

Start with "uncheck all." Then turn on only what the agent needs.

7 Use expert skills for expert work

If you're not an expert, you can't tell whether the insight is good.

For general tasks, the built-in skills are fine. For domain work like analyzing customer feedback, pricing or sales, use skills built by people who know the field. A general skill will happily blend free-tier and enterprise feedback together, which an expert would never do. Lenny's open-source skills repo is a strong place to start: github.com/RefoundAI/lenny-skills

Worried about customer data? Strip names and identifiers from transcripts first with a tool like Microsoft Presidio. The themes stay; the personal details don't.

8 Example: an end-of-day follow-up agent

We started this one live with the ChatGPT group.

It reads the day's call transcripts (from Granola or a Google Drive folder) and drafts follow-ups in Gmail in your voice. Things worth building in:

- Sort each meeting into draft, skip or internal, so it doesn't write a recap of your dentist appointment.
- Check the meeting actually happened before thanking anyone for their time.
- Decide what to do when there's no transcript.
- Track every "I'll send you…" so nothing you promised gets lost.
- Check later meetings that day, in case a decision changed.
- Pull out feedback, themes, action items, owners and dates.
- Send you one short summary (in Slack, if that's where you live) with links to each draft.

Your voice skill makes or breaks this. Every time a draft sounds off, tell it exactly what you'd never say, and it gets a little better.

9 Start with one skill, then split

Ask the AI directly: is this better as an agent or as a simple workflow?

Begin with a single skill. When the output stops improving, break the work into separate skills and agents. For competitive intel, that might be one agent for news scanning, one for competitor product updates and one for positioning and messaging. Another idea from the room: an events agent that researches who's attending a conference and builds your hit list ahead of time.

NEXT STEP

Pick one job you do every week and build the skill for it tomorrow.

Get the files

momentumproductco.com > Services > AI Workshops. The deck, the ChatGPT agent pack, the Claude agent files and a sample folder of simulated invoices and quotes to practise with. Questions? heather@momentumproductco.com

Join the Slack community

300+ people learning AI at every level. Click to join on Slack