

We AI Field Notes

October 23 2026

Build something people can use

From the Signal Room

Hi everyone,

Friday's conversation moved between personal AI assistants, enterprise adoption and how quickly we can build something useful in ChatGPT. We looked at a live interactive example and the possibilities beyond another written answer: a calculator someone can adjust, data they can explore, or an explanation they can work through.

What interests me is the speed at which we can get to a first version. That gives us more time to ask whether it helps someone make a decision. The first version still needs our judgment, a check of its assumptions and some design work.

These notes share the learning without identifying participants or attributing their comments.

Try this: three useful formats

These are three practical ways to approach a build, not three official product names. What appears in your account depends on the experience and features available to you.

1. A calculator for a decision

Start with a question you already work through with a customer or colleague. A calculator makes the inputs visible and lets the person see what changes the result. Useful examples include a workflow-cost estimate, capacity planner or comparison of two options.

Copy and adapt:

Build a working workflow-cost calculator directly in this conversation. Let me change monthly ticket volume, the percentage eligible for the workflow, minutes per ticket and hourly labor cost. Show the inputs, units and formulas. Use clearly labeled example values, not claims about my company. Distinguish modeled labor value from realized cash savings. Include an explanation of what the calculation leaves out. Do not connect to business systems or publish anything.

Check it: Set volume to 10,000, eligibility to 40%, handling time to 8 minutes and hourly cost to \$50. The model should show about 533.3 hours and \$26,666.67 in modeled labor value. Reducing eligibility to 20% should halve both results. Those are illustrative calculations, not a savings forecast. Implementation cost, quality checks and adoption still matter.

2. Interactive data

A chart becomes more useful when someone can inspect the assumptions or focus on the part that matters to them. Start with a small, approved dataset. Ask for the source, missing values and definitions alongside the analysis rather than letting a polished visual hide an uncertain answer.

Copy and adapt:

Use only this approved, anonymized dataset. Create an interactive view that helps me compare [the decision or metric] by [the relevant category or period]. Keep the original numbers unchanged. Identify missing values and explain any calculation. Do not invent data to fill gaps. Show me the main finding and one limitation before suggesting a decision.

Check it: Pick three rows and compare them with the original. Check totals, units and date ranges. Make sure a filter changes the view without changing the underlying definition.

3. An interactive explanation or simulation

Sometimes the problem is understanding a relationship, not producing a number. A small simulation can help someone see what happens when an assumption changes. It can also help a team discuss a process before trying it in production.

Copy and adapt:

Create a simple interactive explanation of [process or decision]. Let me change [one or two important variables] and see the effect. Label assumptions and show what remains constant. Use fictional data. Explain where this simplified model stops being reliable. Keep it in this conversation; do not publish or take external actions.

Check it: Try an ordinary case, an edge case and a missing input. Ask which parts are calculations and which are illustrative. A demonstration is not evidence that a real workflow is connected or running.

For current capabilities and examples, see [ChatGPT visualizations](#) and [data analysis in ChatGPT](#). These exercises extend Friday's discussion; they are not a claim that we completed all three builds together.

Giving a personal assistant useful context

We compared experiences with personal agents and AI Chief of Staff workflows. The practical questions were familiar: does the assistant understand the work, know what matters and follow through?

A useful setup separates three responsibilities:

- **The assistant** helps prepare and carry out approved work.
- **The Chief of Staff role** connects priorities, customer commitments and decisions that need your attention.
- **The operating system** holds the records, instructions and permissions that keep the work consistent.

One tool may do more than one job. The important part is knowing where the current record lives and what requires your approval. A drafted email is not a sent email. An invitation sent is not a meeting accepted.

A small test for your own setup: Give it one recurring task with a clear finish line. Name the sources it may use, the action it must not take without approval and the evidence you want back. Try a missing source or a conflicting date before trusting it with more.

The [Chief of Staff guide](#) includes setup and review guidance. The [public starter-skills folder](#) contains reusable instructions. Compare these with what you already have. Keep stronger permissions and working records; do not replace a useful setup simply because there is a new version.

Claude Duet: announcement versus access

Claude Duet was part of Friday's news discussion. Anthropic announced its proactive-agent research preview on October 21, with access starting through a waitlist for US Pro and Max subscribers. Its announcement describes connected-app briefs and continued task work. That is a vendor description, not a promise that every account has access or that a workflow is ready for your organization. [Read Anthropic's announcement](#).

Model choice is only part of the workflow

OpenRouter came up as something to explore for working with multiple models. Its documentation describes a common API for model access and routing, with tools for building agents. Choosing a model still does not settle which business data it may read, who can approve actions or how you will test the result. [OpenRouter overview](#).

A useful comparison starts with one task and the same approved input. Look at output quality, cost, speed and the checks required before using the result. Keep confidential information out of a new service until its use is approved.

Choosing a first workflow

We discussed the difference between a capability shown in a demonstration and what people can use in their own organizations. Licensing, integrations, access to maintained business information and security ownership all affect the answer.

Before choosing a first workflow, ask:

1. What decision or task should improve?
2. Who owns the source information and keeps it current?
3. Which systems can the agent access, and with what permissions?
4. What needs human review before it reaches a customer or changes a record?
5. What result would show that the workflow is useful?

A written follow-up on where enterprise AI is gaining meaningful traction is still to come. Vertical AI tools for smaller businesses also came up as a possible future topic.

Session materials

- [October 23 session deck](#).
- [Optional Chief of Staff practice prompts](#). These are the prepared fictional simulation, not a claim that it ran during Friday's live session.
- [Optional fictional participant exercise](#).
- [Chief of Staff guide](#).
- [Public starter-skills folder](#).

Coming October 30

We're confirming timing with Amanda for October 30. Amanda Machan of nCino's planned topic is **PLOPs Rewired and Workforce Acceleration**. Check the session page for current details. [Signal Room details](#).

If you try one of the exercises, bring the result and the question it raised. A useful first version gives us something concrete to learn from.

Best.

Lara