

PRODUCTIVITY · 9 min read · September 15, 2026 · By Kelly Stratton

How to Stop Running Out of ChatGPT-6 Astra Usage: 2 Settings and 5 Prompts

Get more done with ChatGPT-6 Astra on the \$20 plan without hitting the limit: the reasoning level to set, how Goal mode keeps Astra on one job, and five copy-and-paste prompts built to finish a whole task in one run.

You paid for ChatGPT-6 Astra, and it can run a whole job for you: sort the inbox, build the spreadsheet, research the purchase, prep the week. Then halfway through the afternoon it tells you that you've hit your limit and to try again in a few hours. That's fixable. Two settings and a better way of asking get you far more out of the same \$20 plan, and you'll stop watching the meter.

IN A NUTSHELL

Keep ChatGPT-6 Astra on medium reasoning, not high. Give it the whole job in one prompt, in three parts: an action verb, the constraints, and what done looks like. For a bigger job on the desktop app, type /goal first so Astra stays on one goal until it's finished. Below: why Astra runs out, exactly where the settings are, and five prompts built to finish in one run.

WHAT "USAGE" MEANS

ChatGPT-6 Astra usage is the allowance of Astra messages your plan gives you inside ChatGPT Work and Codex. On ChatGPT Plus, OpenAI lists 5 to 45 messages every five hours, plus a weekly cap. Both Work and Codex draw from the same pool, and heavier messages use more of it. Regular chat has its own limits.

WHAT YOU NEED

A ChatGPT Plus plan (\$20 a month) or higher. ChatGPT-6 Astra lives inside ChatGPT Work: on your phone, tap the dropdown at the top that says Chat and switch it to Work; on the desktop app, open the ChatGPT menu in the top-left corner and flip Chat to Work. Goal mode needs the desktop app. Everything else in this guide works on the phone too.

PART 1 Why Does ChatGPT-6 Astra Run Out of Usage So Fast?

ChatGPT-5.6 Sol answered questions. You typed, it replied, one message, done. ChatGPT-6 Astra does jobs. It opens a browser, reads pages, clicks, builds documents, checks its own work. Every step it takes, and every question it stops to ask you, counts against the same five-hour allowance. That's why the same \$20 plan feels smaller with Astra: you're paying per job, and a job can be one message or ten.

Three things turn one job into ten messages:

What burns it

Reasoning set too high

Higher reasoning makes Astra think longer on every step, and OpenAI's own guidance says higher effort can use more of your allowance and does not always produce a better result. On most jobs you pay more for the same answer.

Not enough context

Astra will stop and ask when it doesn't know what you want. Each question is a message from it, and each answer is a message from you. Three rounds of "did you mean" and the job has cost four times what it should have.

Wandering

With no clear finish line, Astra keeps going: one more source, one more pass, a rewrite you didn't ask for. It doesn't know it's done because you never said what done looks like.

The fix for the first one is a setting. The fix for the other two is how you write the prompt, and on the desktop app, Goal mode.

PART 2 Which Reasoning Level Should I Use for ChatGPT-6 Astra?

Medium. Unless you're on the \$200 Pro plan, keep ChatGPT-6 Astra on medium and leave it there. Simple prompt, complex prompt, doesn't matter.

Here's why that isn't a compromise. On September 6, an OpenAI staff member wrote in OpenAI's developer forum that GPT-6 Astra on low already performs better than GPT-5.6 Sol on high, and told people who were happy with Sol on high to move down to low or medium on Astra. Medium is already a step above the best setting you had before. Independent benchmarks published the same week found that going from medium to high on Astra bought about one point on their scale for about 22% more cost. You would never notice the point. You will notice the usage.

The five levels in ChatGPT Work

Light

Fastest, uses the least of your allowance. Fine for quick jobs where the steps are obvious: sort these, list those, pull this into a table. OpenAI describes the low setting as the one that makes your allowance go further, and its staff say Astra on low already beats the old Sol on high.

Medium

The default for everything. Enough thinking for multi-step jobs, research, and documents, without the bill. This is where ChatGPT-6 Astra should sit on a \$20 plan.

High

For a job where Medium visibly got it wrong, and only then. On Plus, it's the fastest way to hit the limit.

Extra High and Max

Pro-plan territory, where the allowance is five to twenty times larger. On a \$20 plan, these use the most allowance per step of any setting. Leave them alone.

Where the setting is. Tap the model name at the top of the chat. The reasoning level sits with it, and the same five labels show up on every model: Light, Medium, High, Extra High, Max. The advice in this guide is for ChatGPT-6 Astra in Work, because that's where the allowance is small and the jobs are long. Switch to Work, check that the level says Medium, and leave it there.

PART 3 What Does /goal Do in ChatGPT?

Goal mode keeps ChatGPT-6 Astra pointed at one finish line until it gets there. Instead of doing a step and handing the conversation back to you, Astra keeps working, checks its own progress against the goal you gave it, and stops when the goal is met. That's what stops the wandering, and wandering is what burns your usage.

How to start it. In the ChatGPT desktop app, type `/goal` in the message box, then describe the goal. Three things belong in that description, and they're the same three parts as any good Astra prompt: the outcome you want, the constraints, and what a finished result looks like. Goal mode needs that last one most. If you can't say what done looks like, Astra can't tell when to stop.

While it runs. A progress row appears above the message box. From there you can pause, resume, edit the goal, or clear it. You can also keep typing: send a follow-up to add a detail or tighten a constraint, and Astra folds it in without starting over. If you want to know where it's at, ask it to summarize progress. Keep related work in the same chat so it doesn't lose the thread.

On the phone or in a browser. There's no `/goal` there. Use ChatGPT Work and put the outcome, the constraints, and the finished result directly in your prompt. You get most of the benefit, because the reason Goal mode works is that it forces those three things to be said. Prompt 5 below is written that way.

When to use it. Goal mode is for jobs with several steps and a clear end: build the whole spreadsheet, research all fifteen, get the document to final. For a two-minute job, a normal prompt is enough.

YOUR SETTINGS LIST

1. Jobs go in ChatGPT Work. Quick questions go in regular chat, which has its own limits.
2. Reasoning level while you're in Work with ChatGPT-6 Astra: Medium, and leave it.
3. If you also use Codex, remember it draws from the same Astra pool, so don't run the same job in both.
4. Big job on the desktop: type `/goal` first.
5. One job per chat, and keep the follow-ups in that chat instead of starting a new one.
6. Every prompt ends with "Done looks like:" so ChatGPT-6 Astra knows when to stop.

PART 4 How Do I Write a Prompt That Finishes in One Run?

Three parts, in order. **Start with an action verb:** open, search, read, build. **Give it constraints:** what counts, what to skip, what not to touch, and how to handle a question without asking you. **End with what done looks like:** the exact shape you want back. The full formula, with ten prompts for everyday jobs, is in [How to Prompt ChatGPT-6 Astra](#). The five below add one thing that guide doesn't: each one is built to cost one message.

The line that does most of the saving is this one, and it's in every prompt: *"Don't stop to ask me questions. If something is unclear, make a sensible choice, keep going, and list the choices you made at the end."* That single sentence turns a four-message job into a one-message job.

About the brackets. Anything in [square brackets] is a blank for you to fill in before you paste. Type your real detail in its place and delete the brackets. When Astra reaches a page that needs a login, it shows you a sign-in form. Type your password into that form, not into the chat. The form sends it straight to the browser, and the model never sees it.

1. The one-run inbox

Sort, unsubscribe, and draft the replies, all in one pass. Most people run this as three separate prompts. It's one.

PROMPT

Open my email and go through everything from the last 7 days. Unsubscribe from newsletters that are sitting unread from the last 30 days. Archive receipts, shipping notices, and automatic notifications. Leave everything else where it is. Don't delete anything and don't send anything. For each message that needs a reply from me, write a two-sentence draft in a plain, friendly tone. Don't stop to ask me questions. If something is unclear, make a sensible choice, keep going, and list the choices you made at the end. Done looks like: one list with the sender, the subject, and the draft reply for each one, followed by a count of what you unsubscribed and archived.

2. The buy-or-don't brief

Everything you'd spend an evening reading, in one table with a pick at the bottom.

PROMPT

Research [the product or service you're deciding on] for [what you'll use it for]. Compare the 3 best options for someone who [your situation: a small business, a family of four, a beginner]. Use the companies' own sites and two independent reviews each. Ignore sponsored lists. Don't add anything to a cart, sign up for anything, or contact anyone. Don't stop to ask me questions. If something is unclear, make a sensible choice, keep going, and list the choices you made at the end. Done looks like: a table with price, what's included, the best thing about each, the one drawback, and the return policy, then your pick in two sentences.

3. The long document, decided

Attach a contract, a policy, a report, or a proposal, and get back the version you can act on.

PROMPT

Read the attached document all the way through. Write a one-page plain-English summary for someone who hasn't read it. Then list every date, deadline, fee, and obligation it puts on me, with the page or section it comes from. Then list the 5 questions I should ask before I sign or agree. Don't summarize anything you didn't read, and don't guess at numbers. Don't stop to ask me questions. If something is unclear, make a sensible choice, keep going, and list the choices you made at the end. Done looks like: those three sections, in that order, on one page.

4. The application pack

A posting and your resume go in. A tailored resume, a cover letter, and an honest gap list come out, in one run.

PROMPT

Open this job posting: [link]. Read the attached resume. Rewrite my resume bullets to match the posting's wording where it's true, and don't invent anything I didn't do. Then write a 150-word cover letter in a plain, confident voice that names the company and the role. Then list every requirement in the posting that my resume doesn't show, so I can decide whether to address it. Don't apply, don't submit anything, and don't message anyone. Don't stop to ask me questions. If something is unclear, make a sensible choice, keep going, and list the choices you made at the end. Done looks like: the revised resume, the cover letter, and the gap list.

5. The Goal-mode job

A real multi-step job with a finish line built in. On the desktop app, type `/goal` and paste it. On the phone, paste it into Work as is.

PROMPT

Build a spreadsheet of 15 [type of business or vendor] within 20 miles of [city]. Use each one's own website and their Google Maps listing. For each, get the name, the address, the phone number, the website, their hours, their starting price or price range if it's listed, and their Google rating. Skip any that have been closed or have no website. Don't fill out forms, book anything, or contact anyone. Don't stop to ask me questions. If something is unclear, make a sensible choice, keep going, and list the choices you made at the end. Done looks like: a spreadsheet with those 8 columns and exactly 15 rows, sorted by rating, highest first. Stop when there are 15 rows.

WHY EVERY PROMPT SAYS WHAT DONE LOOKS LIKE

That last sentence is the finish line. Without it, ChatGPT-6 Astra decides for itself when the job is over, and it tends to keep going. With it, Astra checks its own work against your sentence and stops. In Goal mode that sentence is the whole point. In a normal prompt it's the difference between one message and three.

PART 5 The Tell

When a job is eating usage, what comes back tells you which fix you skipped.

What you see, and what to change

It asks you two or three questions before it starts

Your constraints are thin. Add what counts, what to skip, and the "don't stop to ask me questions" line.

It finishes, then keeps adding to the result

No finish line. End the prompt with "Done looks like:" and the exact shape. On the desktop, run it under /goal.

A simple job took a long time and you hit the limit early

Reasoning is on High or above. Tap the model name and set it to Medium.

It did half the job and stopped

It hit something it couldn't get past: a login, a "prove you're human" box, a pop-up. Take over the mouse, finish that one step, and type "keep going." It picks up where it left off without starting the job over.

It says you've hit your limit in the middle of a job

Don't start a new chat when the window opens again. Go back to the same chat and type "keep going from where you stopped." The work it already did is still there, and finishing costs one message instead of the whole job again.

PART 6 What the Settings Don't Change

The limit is still the limit. Medium and one-run prompts stretch the allowance; they don't remove it. The five-hour window resets on its own, and there's a weekly cap on top. ChatGPT tells you when you've hit the window and when it opens again. The honest way to plan is to save Astra for jobs that take clicking and give quick questions to regular chat, which has its own limits.

Codex comes out of the same pool. If you also use Codex, its jobs and your Work jobs share one allowance, so a heavy afternoon in one leaves less for the other.

You can buy a reset, in some countries. On eligible personal Plus and Pro accounts, OpenAI sells an instant reset that refreshes the five-hour and weekly allowance the moment checkout goes through. It's there if you need it on a deadline. The two settings above mean you mostly won't.

DO THIS FIRST

Open ChatGPT Work, tap the model name, and set the reasoning level to Medium. That's the whole first fix and it takes ten seconds. Then paste prompt 2 with something you've been meaning to buy. Watch it come back in one message with the table and the pick. From then on, every job you hand ChatGPT-6 Astra gets the same three parts and the same last line, and your usage goes a lot further.

Straight answers

How many ChatGPT-6 Astra messages do you get on the \$20 plan?

OpenAI's help center lists 5 to 45 Astra messages per five-hour window on ChatGPT Plus, plus a weekly cap. The range is wide because heavier messages, and higher reasoning levels, use more of the allowance. ChatGPT Work and Codex draw from the same pool.

Which reasoning level should I use for ChatGPT-6 Astra?

Medium, unless you're on the \$200 Pro plan. OpenAI says higher effort can use more of your allowance and does not always produce a better result, and OpenAI staff have said Astra on low already performs better than ChatGPT-5.6 Sol on high. Go up only when medium visibly fails.

What is Goal mode in ChatGPT?

A mode that keeps ChatGPT-6 Astra working toward one goal until it's finished instead of stopping after each reply. In the ChatGPT desktop app you start it by typing `/goal` in the message box. A progress row above the box lets you pause, resume, edit, or clear the goal while it runs.

Does `/goal` work on the phone?

No. The `/goal` command is in the ChatGPT desktop app. On the phone or in a browser, use ChatGPT Work and put the outcome, the constraints, and what a finished result looks like directly in your prompt. That gives Astra the same three things Goal mode asks for.

Why does ChatGPT-6 Astra run out of usage faster than ChatGPT-5.6 Sol?

Sol answered a question in one reply. Astra does a job, and every step it takes and every question it asks you counts against the same allowance. A vague prompt makes it ask questions, wander, and start over, so one job can cost several messages. A three-part prompt with clear constraints usually costs one.

Can I buy more ChatGPT-6 Astra usage?

On eligible personal Plus and Pro accounts, OpenAI sells an instant reset that refreshes your five-hour and weekly Work and Codex allowance as soon as checkout goes through. Availability depends on your billing country. Otherwise, the five-hour window resets on its own.

Do ChatGPT Work and Codex share the same limit?

Yes. OpenAI's release notes say ChatGPT Work and Codex share the same usage limits and credits. A job you run in Codex comes out of the same pool as a job you run in Work, so plan both together.

SOURCES Where These Numbers Come From

- OpenAI Help Center, "Managing usage with GPT-6 Astra in Work and Codex": the 5 to 45 messages per five-hour window on Plus, the weekly window, the effort guidance, and purchased resets. help.openai.com
- ChatGPT Learn, "Long-running work": the /goal command, the progress row, and the web guidance. learn.chatgpt.com
- ChatGPT Learn, "What's new," August 3 to 7, 2026: Work and Codex share the same usage limits and credits. learn.chatgpt.com
- OpenAI Developer Community, September 6 to 7, 2026: OpenAI staff on Astra low versus Sol high. community.openai.com

Want *more* like this?

Free AI guides added regularly. No jargon, no overwhelm.

kellystrattonai.com

Instagram / TikTok: @kellystratton.ai | Facebook / YouTube: @kellystrattonai