PUBLIC SERVICES ALLIANCE

GPT-6 Astra

Your practical reference for ChatGPT, Work, and PSA / IHPF

Settings explained • Recommended choices • Verified changes
Useful functions • Ready-to-use PSA workflows

TLDR — what to do now

Use **Work** for a finished report, spreadsheet, presentation, or other substantial deliverable. My recommended starting point for demanding IHPF work is **Astra Medium, Standard speed**, if available. Use a lighter model for routine editing and extraction.

Keep PSA work in the PSA Business workspace, provide the current source documents, and state what a usable result must contain. You do not need programming knowledge to use these capabilities.

Verified against official OpenAI documentation on September 5, 2026.

This is a comprehensive practical guide to documented capabilities. Account-specific availability and control labels can differ. Recommendations are advisory; no settings or subscriptions were changed.

Official sources: [1] · [2] · [34]

How to use this reference

Read the quick start first. Then use the settings pages when configuring an account, the changes pages when comparing generations, and the workflow pages when assigning work. Quick start

Names and model selection

Reasoning and speed

Personal and desktop settings

1. Start here: your first ten minutes

PSA recommendation: begin with one useful, bounded IHPF deliverable and learn the controls while producing it.

1. **Check the workspace.** Confirm that the account/workspace selector shows PSA Business before adding organizational material.
2. **Open the intended IHPF project.** Start the new chat from inside it. Confirm that it appears in that project's Chats list.
3. **Choose the experience.** Use Chat for a quick explanation; select Work for a completed deliverable.
4. **Choose the model.** In Work, inspect the control beneath the prompt box. Open Advanced if offered. Select GPT-6 Astra explicitly when you want to test Astra.
5. **Start at Medium and Standard speed.** These are my recommended trial choices, not a claim about your account's defaults.
6. **Attach the governing material.** Provide the latest approved IHPF document and any competing proposal. Identify which version governs.
7. **Name the output and checks.** Ask for a Word decision brief, citations, assumptions, unresolved questions, and readable formatting.
8. **Review the result.** Open the file, inspect its evidence and layout, and request specific changes in the same chat.

Create a board decision brief comparing these two IHPF proposals. Identify the decision, evidence, costs, partner dependencies, and next steps. Label assumptions.

Official sources: [4] · [8] · [34]

2. The names: what each one means

Term	Meaning
GPT-6 Astra	The model doing the reasoning and writing.
ChatGPT	The application through which you work with models and tools.
Chat	An experience for questions, discussion, and short drafts.
Work	An experience for carrying a substantial task to a reviewable result.
Codex	An experience that exposes more software-development detail.
Business workspace	PSA's organizational account boundary and administration.
Project	Related chats, instructions, and sources organized around a goal.
Plugin / connected app	Additional workflows, information, or supported actions.
API	A developer's way of using models in separate software.

PSA recommendation: use Chat and Work for everyday nonprofit work. Open Codex when building or maintaining software requires its technical views. Selecting Work does not itself establish which model is selected.

Cloud versus local: cloud work uses hosted resources. Local work uses authorized resources on your Mac. A cloud task cannot automatically see Finder files or the browser tabs open on your computer.

Official sources: [4] · [15]

3. Select and identify the model

PSA recommendation: select the named model for an important comparison or reproducible workflow. Record the visible name and reasoning choice with the result.

Where to look

Work on web or desktop: inspect the model/reasoning control below the composer—the box where you type. **Advanced**, when offered, exposes more specific choices. Ordinary Chat can have a different picker; do not transfer a Work menu path to it.

If the label says “Latest”

Treat it as an unresolved selector label until the app exposes its meaning. The documentation reviewed does not establish a stable “Latest = Astra” rule for your account. Open the selector and inspect the named choices. If no exact name is shown, retain a screenshot of the menu and the app version.

Asking the assistant “What model are you?” is not a reliable audit method. This guide confirms that Astra is documented; it does not certify the model behind every message in your account.

For this PSA task	My suggested model
Routine extraction or a fixed-format list	Luna
Ordinary editing, summaries, and simple reports	Terra
Substantial synthesis and polished writing	Sol
Difficult research, complex decisions, or a complete multi-tool deliverable	Astra

If Astra is missing, check the signed-in workspace, product experience, rollout, and admin access. A personal setting cannot grant a model your account is not eligible to use.

Official sources: [2] · [12]

4. Reasoning, speed, Max, and Ultra

Control	PSA recommendation
Light / Low	Use for a clear, short task.
Medium	Start here for substantive IHPF work.
High / Extra High	Use for conflicting evidence, complex budgets, or major strategy choices.
Max	Reserve for a difficult problem that a lower effort did not resolve.
Ultra	Use deliberately for a large task with genuinely independent parts.
Standard speed	Keep as the normal choice.
Fast mode	Use when reduced waiting is worth higher usage.

Documented distinction: reasoning effort affects how much analysis the model undertakes. Max emphasizes one task's depth. Ultra can distribute work across subagents. Fast mode concerns processing speed and has a higher credit rate; Astra Fast is documented at 2.5 times the Standard Work/Codex rate.

Where: use the Work model/power control and Advanced, if offered. Desktop documentation places the optional Ultra slider control under Settings → Configuration. Do not assume that every effort level appears on every client.

My trial rule: give Astra Medium three representative PSA tasks. Increase effort when the result misses an important tradeoff or consistency check. Improve the prompt or sources before assuming every failure requires Max.

A higher setting does not create missing evidence, grant access to files, or guarantee correctness. Asking for current sources is a separate instruction.

Official sources: [2] · [11]

5. Personalization: set the working style

These recommendations apply to the person using the account. Organizational requirements also belong in the relevant project instructions.

Setting	My recommendation
Personality	Pragmatic in the documented desktop menu; Candid if that is the option your Chat menu actually offers.
Custom instructions	Add a short, explicit PSA working standard.
Memories	Enable for useful continuity where available; review the retained context.
Computer History	Leave off initially; assess separately if recovering recent Mac activity becomes valuable.

Where: Settings → Personalization for personality, instructions, and memory controls. The desktop documentation lists Friendly, Pragmatic, and None. Personality affects communication style, not model capability.

Suggested personal instructions

Lead with a brief TLDR. Be candid, professional, and direct. Explain unfamiliar terms in plain English. For settings, name the product and platform, give the documented location, and say when a control may be unavailable. Do not invent menu items. For documents, use 13-point body text, no text under 12 points, clear headings, and grids around tables. Never reduce type to meet a page limit.

Memory is supporting context. Web memory and local Codex memory have different stores and controls. Keep the approved PSA mission, current pilot scope, and document standards in project instructions and source files so colleagues do not depend on your personal recall.

Official sources: [6] · [7] · [18]

6. Desktop comfort and accessibility

Open desktop Settings: use the app menu, or press **Command + comma** on your Mac. These locations are documented for the current desktop app, not for every ChatGPT web screen.

Setting / location	My recommendation
General → multiline submission	Require Command + Return/Enter if long prompts are being sent accidentally.
General → Follow-up behavior	Steer the active run when your next message should correct current work.
Appearance	Choose high contrast and comfortable UI text. Prefer a quiet theme.
Keyboard Shortcuts	Search by command name; inspect the actual assigned keys.
Notifications	Notify while the app is in the background; enable questions and approval alerts.
Suggested prompts	Turn off if they add clutter; keep only if they help you resume work.
Pets	Optional; leave off unless you find the activity indicator useful.

Keep the prompt accessible

The documented desktop option is to **pop out a chat** and enable **Always on top**. The documentation does not establish an exact icon location for your version. Use the available chat/window menu; if that option is absent, do not keep hunting for a supposed hidden drag bar.

Practical fallback: keep the normal ChatGPT window beside your work. Enlarge text with the Zoom menu you previously found, or Safari's View menu. These workflows do not require Quick Chat.

Official sources: [5] · [28]

7. Work controls: completion and access

Control / location	My recommendation
Desktop composer → Work locally / Cloud	Cloud for research and uploaded files; local only when Mac files or apps are necessary.
General → Prevent sleep while running	On for local jobs. It is unnecessary for a hosted task's continued execution.
Settings → Cloud browser	Auto approve for ordinary public research if offered; Always ask for sensitive work.
Desktop Settings → Browser	Allow needed sites; review access when the workflow changes.
Data controls → Work network access	Keep public command networking off unless a task needs it.
Local file/app permissions	Grant access to the folders and apps needed for the assigned task.

Why the network distinction matters: the Work network toggle controls code/shell internet access. It is separate from search, connected apps, and the browser. Turning it off does not mean all external access is disabled.

For long tasks: tell Work the intended result and the checks that define completion. Use the same chat to correct the direction. In a supported desktop workflow, /goal provides a persistent goal; on the web, put the goal directly in the prompt.

Finish the comparison and create the review copy. Continue through routine reversible steps. Keep me informed about substantive findings. Ask only when missing information would materially change the result or an action requires approval.

Official sources: [15] · [17] · [27] · [34]

8. Projects, sources, memory, and files

PSA recommendation: keep one clearly designated current IHPF reference set, with separate chats for each outcome.

Put the right material in the right place

Project Sources: material needed across the project. **Chat attachment:** material for one request. **Memory:** helpful prior context. **Library:** retained files where available. **Connected Drive/Dropbox:** documents retrieved through an authorized connection. These are different mechanisms.

On the web, open the project's **Sources** section for shared context and start chats from its **Chats** section. A file existing somewhere in your account does not establish that a particular chat has read it.

PSA's recommended source set

9. A dated "Current IHPF direction" document naming the governing strategy and approved scope.
10. A source index with title, date, owner, status, and link to each important resource.
11. A decision log distinguishing approved decisions, proposals, and unresolved questions.
12. A small set of accepted examples: board minutes, research comparison, and public presentation.

Before a major synthesis, ask: **"List the documents you actually opened, their versions, and any source you could not access."** A reference to a file in an old conversation is not proof the full file was available now.

After creation, open or download the actual Word, Excel, PowerPoint, or PDF file. Ask for a filename and a short account of the checks performed. Keep an independent copy of approved PSA records using your established filing process.

Official sources: [8] · [9] · [15]

9. Understand usage before changing plans

PSA recommendation: Standard speed by default; compare task value and observed usage before buying more capacity.

Work and Codex share usage. Actual consumption varies with model, context, tools, reasoning, and task length. A credit draw is not automatically a new invoice charge. The applicable seat, allowance, credits, and overage arrangement matter.

Standard Work/Codex model	Input / output credits per million tokens
Astra	250 / 1,250
Sol	100 / 500
Terra	50 / 300
Luna	5 / 30

These published token rates apply to eligible credit billing, not a fixed price per question. Cached-input rates differ. API dollar pricing is separate. Do not convert credits into dollars without PSA's applicable purchase or contract rate.

What to check in your account

13. Your exact Business seat type, remaining allowance or credits, and next reset date.
14. Whether extra usage stops, consumes prepaid credits, or can create overage charges.
15. Whether a workspace spending limit, individual limit, and alerts are offered.
16. The current billing period and actual invoice, separately from estimated usage.

Published guidance distinguishes Standard Business and Business ($100), and rollout descriptions are not fully uniform across pages. Do not infer that PSA needs a $100 seat—or already has Astra entitlement—from one table. Inspect the actual account before making a purchase.

For your four users, log a few completed tasks: minutes saved, corrections required, usefulness, and reported consumption. This is more informative than counting prompts.

Official sources: [10] · [13]

10. Administrator settings for PSA

Open PSA's administrative or Workspace settings area from the account/workspace menu. Labels and editability depend on your plan and role.

Setting	PSA recommendation
Models / Work & Codex defaults	Terra or Sol for routine team use; permit Astra for substantial work if eligible.
Work access	Enable for the four intended PSA users where supported.
Default speed	Standard.
Plugins / connected apps	Allow the few services PSA actually uses; start with document sources.
Usage limits	Apply a deliberate trial allowance where your account offers an enforceable control.
Workspace overage	Zero or an explicitly approved cap if overage is supported.
Usage alerts	Enable useful thresholds; remember that alerts do not cap spending.
Public sharing / publishing	Require an intentional review before external release.
Computer History	Do not enable broadly as part of a routine model upgrade.

Where the documentation points: model defaults under Workspace settings → Models; credit balance, reports, and applicable alerts under Billing. Source-service permissions remain separate from ChatGPT permissions.

Analytics: the reviewed workspace-analytics guidance concerns Enterprise/Edu coverage. I cannot promise a Business "Workspace analytics" menu. Use the billing/usage reporting actually available to PSA.

Official sources: [12] · [13] · [14] · [36] · [37]

11. Privacy and IHPF boundaries

PSA recommendation: use public research, approved organizational documents, and synthetic examples for the first wave of IHPF workflows.

Business data is not used for training by default under the documented plan guidance. That fact does not mean every plugin, sharing choice, or clinical workflow is automatically appropriate.

Keep four boundaries distinct

17. **Workspace:** select PSA Business for PSA institutional work. Keep your personal work in the intended personal workspace.
18. **Project sharing:** inspect the project's actual audience. An unshared project inside Business is not a shared collaboration space merely because four people have seats.
19. **Connected service:** authorization can expose the connected account's permitted records. A shared connection may have different access from an individual user.
20. **Publication:** a chat, document, and hosted Site can have different audiences. Review the specific item's access settings.

Healthcare data and clinical decisions

Astra is not a blanket authorization to upload identifiable patient records or deploy autonomous diagnosis. OpenAI's HIPAA guidance describes eligible products, agreements, configuration, and third-party responsibilities. PSA's ordinary Business setup should not be assumed to meet those requirements.

For a navigator or screening pilot, use synthetic cases until the relevant clinical partner has approved the data flow, human responsibilities, escalation process, and applicable agreements. Measure whether staff can use the output safely and accurately before expanding its role.

This guide proposes organizational uses and prototypes. It does not establish clinical effectiveness, regulatory clearance, or readiness for patient care.

Official sources: [10] · [15] · [29]

12. What improved in GPT-6 Astra

Documented model changes relative to earlier models, including Sol. The statements below summarize OpenAI's claims; they are not PSA's own measured results.

Reported improvement	What PSA should test
Stronger reasoning, science, and software engineering	Can it reconcile IHPF evidence or build a working prototype with fewer repairs?
Better sustained, multistep execution	Can it finish a research brief, workbook, and decision summary consistently?
Stronger instruction following	Does every page respect the minimum text size and required structure?
Better incorporation of corrections	Does a changed requirement propagate through the whole deliverable?
Improved browser and visual judgment	Can it identify a mismatch between a screenshot and a proposed interface?
Stronger attention to task boundaries and transparency	Does it clearly distinguish completed work, assumptions, and blocked steps?
Greater efficiency in some evaluations	Does the completed PSA task need fewer retries or less overall usage?

My interpretation for PSA: the largest opportunity is giving the model responsibility for a complete, testable work product. Evaluate the whole result: evidence quality, consistency, usability, and time saved.

OpenAI also describes a tendency to ask more clarifying questions, give detailed formatted answers, and test coding changes extensively. Use explicit instructions about autonomy, length, and proportionate checks when these behaviors slow a straightforward task.

No universal accuracy percentage, speed improvement, cost saving, or medical performance guarantee is asserted here. Those would require a specified evaluation and comparable task conditions.

Official sources: [1] · [17] · [30]

13. New model mechanics; existing tools

Newly documented Astra mechanisms

21. **Asynchronous tools:** an application can let the model continue other work while a tool is still running. PSA benefit: a potential reduction in idle time in a future custom research workflow. This requires supporting software; it is not a desktop toggle.
22. **Mid-turn steering:** the API supports adding guidance during a run while preserving completed work. PSA benefit: correcting scope while a task is underway. The visible ChatGPT behavior still depends on the client.
23. **Reasoning updates with cache preservation:** developer software can change reasoning effort without rebuilding the original prompt prefix. PSA benefit: a custom system can spend more effort only when needed.
24. **Misalignment monitoring:** OpenAI documents background monitoring and possible interventions. Treat any resulting pause as a system notice, not evidence that the requested deliverable is complete.

Do not label these “invented by GPT-6”

Web search, computer use, structured outputs, image-generation tools, skills, multiagent orchestration, and document workflows already existed. Astra can improve how those capabilities are used; their availability is a separate product question.

Technical specifications — for context only

The API model page lists a **1,050,000-token context window**, up to **128,000 output tokens**, and an **April 30, 2026 knowledge cutoff**. It accepts text and images and returns text; audio/video are not native modalities on that model endpoint. Other tools supply voice and image creation.

These API specifications do not promise an identical ChatGPT file limit or unlimited recall. Experimental context management is separately documented and was not available to Business at launch.

Official sources: [1] · [2] · [3] · [31] · [32] · [33]

14. Recent application changes to know

These are recent ChatGPT/Work developments, not an exhaustive list of model changes. Check account availability before making a workflow depend on them.

Recent capability	Potential PSA value
Event-triggered scheduled tasks	Prepare a brief when a supported message or repository event arrives.
More Chromium browser integrations	Use a supported existing browser tab for a bounded task.
Cloud browser sign-in on eligible accounts	Complete a task requiring a website login when the account supports it.
Site co-editing and editable Site URLs	Collaborate on an internal PSA resource portal.
Computer History	Find the document or activity you were working on earlier.
Local thread snapshots	Share a read-only account of a completed workflow.
Voice inside an existing task	Discuss progress or redirect work verbally.

The official weekly digest highlights Astra for August 31–September 4 and several application changes in the preceding weeks. The exact release boundaries matter: an older capability may be newly available to your account without being a GPT-6 invention.

PSA priority: finished files, connected research, reusable workflows, and useful visual tools deserve attention first. Browser extensions and Computer History are optional conveniences. They do not need to be enabled to get value from Astra.

Safari note: the browser-extension list reviewed names Chrome, Edge, Brave, Opera, and Vivaldi. Do not assume that a Safari extension exists. You can continue using ChatGPT in Safari without adopting a new browser integration.

Official sources: [18] · [19] · [20] · [23] · [30]

15. Research and spreadsheet capabilities

Current-source research

Ask explicitly for web research and direct source links when dates, policy, availability, or product claims matter. Work can combine relevant files and retrieved material into a comparison or brief. A higher reasoning choice alone does not demonstrate that a search occurred.

Research current consumer-facing AI healthcare services. Separate announced, pilot, and generally available offerings. Cite primary sources, record the verification date, and compare them against the attached current IHPF strategy. Mark unavailable evidence.

PSA use: maintain an evidence register for Amazon Health AI and alternatives. Ask for the service's users, geography, access conditions, human support, data requirements, and evidence of actual deployment.

Data analysis and Excel workbooks

Provide a spreadsheet or dataset and request cleaning, formulas, comparisons, charts, and explicit assumptions. Define units and the intended decision. Ask for an editable workbook when others must change inputs.

Create a pilot-cost workbook with separate assumptions, calculations, scenarios, and summary sheets. Label invented estimates as assumptions. Show cost per participant, staffing hours, and sensitivity to attendance. Check the arithmetic and explain missing inputs.

My review standard: change two inputs and confirm that totals update sensibly. Require sources for factual inputs and an assumptions log for estimates. A plausible looking chart is insufficient.

Deep Research label: if your Chat interface offers a separate Deep Research tool, treat it as its own workflow. The pages reviewed do not verify its exact current menu location or GPT-6 mapping for PSA. A Work research task can still be specified directly.

Official sources: [9] · [16] · [34]

16. Documents, slides, and visual tools

Output	What PSA should request
Word reference or board brief	Editable text, readable headings, citations, stable terminology, and a checked layout.
PowerPoint presentation	A clear audience and decision, speaker notes, source notes, and the supplied PSA logo/template.
PDF	A checked printable version when the layout should stay fixed.
Interactive visualization	A small calculator or scenario explorer with transparent assumptions.
Generated image	An explanatory illustration or campaign graphic with a specified style.
Site	A reusable hosted portal or tool with deliberate access and persistent data requirements.

Where: describe the required file in Work. For an interactive tool, type @Visualize and select the suggestion if available. For a hosted tool, ask for a website or invoke Sites. On the web, Sites management is documented under More → Sites.

Use the right medium

For pilot economics, ask for an exact spreadsheet or calculated chart. For community outreach, ask for a simple visual story. For a reusable navigator resource directory, consider a Site. Do not use a generated image as the source of numerical truth.

Give the model a good existing PSA example and identify what must remain fixed. Request a specific revision: “Keep all wording; enlarge the Theme column or replace that table with readable sections.”

A Site is a separate hosted product. Request a saved review version before deployment if you want to inspect it first. Do not assume that an interactive graphic inside a chat is already a durable shared application.

Official sources: [9] · [23] · [24] · [25]

17. Connected apps: start with PSA's files

PSA recommendation: use the connected document sources you already rely on before adding more services.

In this session, Dropbox and Google Drive capabilities are available. That does not prove that every folder or document is accessible to every PSA member. Source permissions and each connection's authorization still determine access.

A simple access test

25. Ask the assistant to find one specifically named, nonsensitive PSA document in the intended service.
26. Require it to open the document and report its title, modification information if available, and a short passage summary.
27. Repeat the test from another PSA user's account before relying on the workflow for the whole team.
28. If it fails, correct the connection or document permissions. Do not treat a search snippet as proof of full access.

Where: use Plugins in the left sidebar where available. Select an installed service or type @ and choose it. Some bundles require an additional connection or sign-in; installation and account authorization are distinct.

What to add later

Email, calendar, or meeting tools may help with agenda preparation and follow-up if PSA already uses those services. Connect one for a concrete workflow. Do not install an entire catalog just because it is available.

Ask for read-and-draft work first. When an action would send, move, delete, publish, or change a record, specify the intended scope. For board communications, have the assistant prepare the message and recipient list for review before sending.

Official sources: [21] · [22] · [37]

18. Reuse work: skills and scheduled tasks

Skills: a repeatable method

A skill packages instructions and supporting material for a recurring task. A plugin can bundle skills and connected tools. Ask the available skill creator to turn a proven PSA workflow into a reusable method; include the accepted template and verification rules.

First three PSA skills I recommend: board minutes with decision/action extraction; current-source healthcare comparisons; and readable PSA document formatting. Test each against a real example before sharing it.

Scheduled tasks: repeat it later

Manage tasks in **Scheduled** where available. A saved task can run on a schedule; eligible web/mobile accounts also support selected Gmail, Slack, and GitHub events. Event triggered tasks are a different trigger type, not continuous universal monitoring.

For PSA's existing daily briefing, first list and inspect the existing task. Check its instructions, destination, timezone, sources, and latest run. Do not create a second briefing merely because you changed models.

A cloud task can work from accessible uploads and connections without your Mac remaining awake. A local task needs its computer and files available. Choose explicitly when setting up the workflow.

Recommended design for recurring outputs

Use a stable prompt with: source scope, time window, maximum number of findings, why each matters, handling of "nothing new," and a clear destination. Specify New Mexico local time and daylight-saving expectations. Review the first few runs before trusting the cadence.

Scheduling a briefing does not automatically share its results with all four PSA users. Test the destination and access separately.

Official sources: [20] · [21]

19. Voice, screen context, and demonstrations

Voice versus dictation

Dictation turns speech into prompt text. **Voice** is a live spoken exchange and can coordinate work. The current desktop voice system uses GPT-Live; the selected task model handles task work. “Voice” does not mean Astra natively processes audio.

Where: Start voice chat or Start new voice chat when shown; Settings → Voice → Voice chat hotkey for a shortcut. My recommendation is to begin with dictation if you want to review your words before sending.

Screen context

On supported Macs, Settings → Voice → Screen context lets you deliberately show the current window. My recommendation is to leave it off until you want help with a visible interface, then show only the relevant window. The captured context can include accessible text beyond the visible scroll area.

Computer History

This opt-in feature summarizes activity across permitted Mac apps and sites. It starts after you enable it; it cannot reconstruct activity it never collected. Business access first requires administrator permission. My recommendation remains off initially.

Record & Replay

For a stable task you already know how to do, desktop Work/Codex → Plugins → + → Record a skill can turn a demonstration into a reusable skill, where Computer Use is enabled. Try it later for a repetitive reporting or filing process.

Record & Replay is a workflow demonstration feature. Do not confuse it with a meeting audio recording/transcription feature. For PSA minutes, a supplied Zoom transcript remains a straightforward source.

Official sources: [18] · [19] · [26]

20. IHPF workflow: evidence to decision

Priority 1 — create a maintained comparison of external AI healthcare services against the current IHPF approach.

Goal: identify where PSA should partner, differentiate, test, or defer. This is a proposed workflow, not a claim that the latest PSA strategy has been reanalyzed in preparing this guide.

Inputs

The most current approved IHPF document, current PSA website posts, relevant prior analyses, and a dated list of external services. Require the assistant to open the sources it actually uses.

Work assignment

Use Astra with sufficient reasoning for the comparison. Build a source-backed matrix of the named healthcare AI services and IHPF. Separate live services from announcements. Compare target users, rural access, human support, pre/post-visit assistance, prevention, diagnostics, costs, evidence, and local community participation. Identify complementary partnerships and direct overlap. Return a Word decision brief and an editable comparison workbook.

What must be checked

29. Each material service claim has a source and verification date.
30. PSA proposals are labeled as proposals; implemented services are not inferred from aspirational language.
31. No exact market share, outcome improvement, or price is invented.
32. Recommendations identify a decision and the evidence that would change it.

Success measure: the board can identify two or three practical next steps and the missing evidence for each. A long catalog without a decision is not enough.

Suggested model: Astra Medium initially; High/Extra High if the source record is large or contradictory. This is my task-specific recommendation.

Official sources: [16] · [34]

21. IHPF workflow: a pilot decision packet

Priority 2 — turn one candidate pilot into a concrete packet partners can assess.

For your discussed micro-lab, ultrasound, cardiac-screening, or eye-screening concepts, start with one use case and one community. Do not ask the model to design all of rural healthcare at once.

Ask for five linked outputs

33. **Service journey:** what happens before, during, and after the encounter, including navigator and clinician roles.
34. **Operating assumptions:** staffing, training, hours, equipment, transportation, and referral capacity.
35. **Budget scenarios:** conservative, middle, and higher-volume assumptions with formulas and clearly labeled estimates.
36. **Partner questions:** clinical responsibilities, data flow, procurement, follow-up, and payment arrangements needing verification.
37. **Evaluation plan:** access, completion of follow-up, staff burden, user experience, and costs; specify how each would be measured.

A useful instruction

Produce a reviewable pilot concept, not a claim of clinical readiness. Mark every assumption and unresolved partner dependency. Distinguish the AI-supported administrative tasks from clinical decisions. Include reasons to continue, revise, or end the pilot, stated in plain language.

Suggested setup: Work + Astra Medium or High + Standard speed. Supply the partner's actual constraints and relevant current sources. Ask for a Word packet and linked budget workbook.

Success measure: a partner can explain the proposed workflow, identify its responsibilities, and name the conditions required for a small test. Quantitative thresholds must be agreed with the partner rather than invented by the model.

Official sources: [9] · [34]

22. IHPF workflow: navigator support

Priority 3 — prototype useful assistance around the clinical encounter using synthetic examples.

Candidate functions

38. Organize a person’s stated concerns into a visit-preparation checklist.
39. Translate approved instructions into plain language while preserving their meaning.
40. Find community resources from a maintained directory with dates and contact details.
41. Draft follow-up reminders from clinician-approved instructions.
42. Help navigators identify incomplete forms, missing information, and unresolved administrative steps.

Build a small prototype

Create five fictional cases covering different language, transportation, digital-access, and follow-up barriers. Ask for an English and Spanish draft, a resource directory, and a staff facing checklist. Have appropriate staff review the language and routing.

Build a navigator-support demonstration for these fictional cases. Use only the supplied approved guidance and resource list. When the information is missing, show the gap and route the question to a named human role. Keep patient-facing language readable and avoid unsupported medical instructions.

Possible output: an interactive demonstration for staff training; later, an access controlled Site if durable records are needed. Define who may see records and how corrections are handled before introducing real information.

Success measure: staff can spot wrong or missing guidance, the tool never invents a resource, and every unresolved issue has a clear human destination. Passing synthetic cases does not establish clinical safety or effectiveness.

Suggested model: Astra for prototype design and evaluation; consider a lighter model only after measuring performance on the approved task.

Official sources: [23] · [24] · [29]

23. PSA board work, outreach, and education

Board meetings and institutional memory

Provide the transcript, agenda, and any formal decisions recorded separately. Request a cleaned transcript, draft minutes, and a decision/action log as distinct outputs. Require “unclear” where a speaker, vote, owner, or deadline cannot be established.

My success criterion: no invented motion, vote, commitment, or attendance. The board reviews minutes before treating them as official. Suggested model: Sol or Astra for messy, ambiguous source material; Terra for a clean routine transcript.

Partner, legislative, and funding outreach

Ask for current-source research followed by a tailored briefing and draft correspondence. Name the recipient role, policy objective, geographic scope, and requested next step. Verify funding deadlines and eligibility against the current issuing source.

My success criterion: every draft uses the correct recipient and current program facts; no email is sent or application submitted simply because a draft was requested.

Public education and EdTech

Turn a vetted PSA explanation into a readable handout, short slide deck, illustration, and interactive scenario. For education research, compare the quality of evidence and implementation conditions rather than equating a product demonstration with proven learning outcomes.

My success criterion: the audience can explain the central idea and the proposed action without needing technical vocabulary. A bilingual version gets appropriate human language review.

A practical four-person division of work

Assign one person to sources, one to draft review, one to numbers and assumptions, and one to the decision/action log. These are proposed responsibilities, not assignments to named colleagues. Let the model prepare the materials that make those reviews faster.

Official sources: [9] · [16] · [25] · [34]

24. Prompt library: stronger first requests

The following are original PSA templates. Replace bracketed material and attach the relevant sources.

A. Comprehensive research reference

Create a comprehensive reference on [topic] for PSA/IHPF. Use the current approved PSA sources and current primary external sources. Begin with the main takeaways. Distinguish verified facts, announcements, assumptions, and recommendations. Include what the evidence cannot establish, practical implications, and source links. Finish the requested file with 13-point body text and no text below 12 points.

B. Exact reformatting

Reformat the attached document without shortening or rewriting its substantive text. Use clear colored headings, 13-point body text, and no text below 12 points, including tables and notes. Use visible grids for necessary tables. Replace an unreadable narrow table with a wider or sequential layout. Allow as many pages as needed. Check every rendered page.

C. Comparison that supports a decision

Compare [options] for [specific decision]. Put my recommended option or setting beside each item. Explain benefits, costs, limitations, implementation requirements, and the evidence. Identify which facts would change your recommendation. Return a concise decision brief and the supporting comparison.

D. Interface help without guessed controls

Use this screenshot and current official documentation to help me change [setting]. Identify the app, platform, and experience first. Tell me one visible next action at a time. Do not name an icon, menu, shortcut, or control unless it is visible or documented for this interface. If your instructions do not match, say so and revise them.

Official sources: [35]

25. Prompt library: direct and verify work

E. Continue without needless pauses

Complete the authorized task. Make reasonable assumptions for routine reversible choices and identify material assumptions in the result. Ask a question only when its answer would materially change the outcome or you need a required approval. Prepare a concrete review copy before asking me to approve an external action.

F. Correct a running task

Keep working on the original deliverable. Apply this correction throughout: [correction]. Preserve the requirements already accepted. Briefly tell me what changed, then continue to completion.

G. Audit evidence and missing access

Check the draft against the original sources. List the files and pages you actually opened. Identify unsupported claims, inconsistent dates, invented numbers, and missing material. Correct what you can verify. Clearly separate the remaining gaps from your conclusions.

H. Finish and report status

Tell me what is completed, what remains, and whether anything is waiting for my input. If the work can continue, continue it. When finished, provide the actual deliverable, explain the checks performed, and identify any material limitations.

I. Delegate only where useful

For this large comparison, split independent source-review tasks among subagents if supported. Give each a defined scope. Reconcile conflicts and produce one consistent final result. Do not have multiple agents edit the same document simultaneously.

Use the last template deliberately for a large parallelizable assignment. It is not needed for a short rewrite or a single document-formatting correction. Additional agents do not guarantee independent evidence.

Official sources: [1] · [27] · [35]

26. Troubleshooting: find the real boundary

Problem	Next useful action
Astra is not listed	Check workspace, experience, client version, and model access. Capture the actual picker.
The picker says Latest	Open its details. If no exact model is exposed, record that uncertainty.
A control is greyed out	Read its explanation; check the role or policy that owns it.
A chat is missing from IHPF	Inspect the project's Chats list and the workspace. Search the actual chat title.
A source is not used	Ask which files were opened; attach or connect the correct version.
A file will not preview	Download and open it in Word, Excel, PowerPoint, or Preview; ask for a fresh output if the file itself fails.
Work seems to have stopped	Inspect progress and pending questions. Ask for completed work, remaining steps, and the exact blocker.
A correction is ignored	Use the same chat; explicitly say the new instruction modifies the active deliverable.

For your prior preview problem: a model upgrade does not prove that preview is fixed. Distinguish a valid downloaded file from a failed in-app preview. Keep the error and a nonsensitive example for support.

For a missing shortcut or menu: do not repeatedly follow an instruction that does not match the interface. A screenshot of the full menu and the app version is more useful than another guessed path.

Official sources: [8] · [9] · [12] · [27]

27. Questions this guide cannot certify

A comprehensive reference is most useful when it makes the remaining unknowns easy to resolve.

Unresolved account question	Evidence needed
What model does PSA's Latest select?	The current selector's details or authoritative account-specific documentation.
Do all four users have the same Astra access?	Each member's visible choices and effective workspace access.
What does an extra task cost PSA?	Applicable allowance, credit rate, billing terms, and measured consumption.
Which greyed-out settings can you change?	The exact setting label, your role, and any policy explanation.
What are the current project/file limits?	The account's current upload/limit notice or applicable current plan guidance.
Is a cloud website login available in Business?	The actual task's offered sign-in capability; the reviewed pages vary in rollout scope.
Does an old interface label still exist?	A current screenshot or release-specific documentation.

I have not invented a new file cap, assumed that a missing analytics menu now exists, or treated API limits as ChatGPT limits. Nor have I assumed the latest IHPF source documents were opened merely because their names appeared in prior conversation context.

Useful next evidence: screenshots of the model picker, Settings categories, and PSA's Billing/usage overview, with private information obscured. They would allow an account-specific companion checklist without rewriting the conceptual guide.

Official sources: [10] · [12] · [14] · [17]

28. Recommended first-month adoption

Use the upgrade to improve a few recurring PSA outcomes and judge the completed work.

Week 1 — establish a usable baseline

Choose the governing IHPF sources. Apply the readable-document instructions. Test one research comparison, one meeting packet, and one budget task. Record model, reasoning, speed, time, corrections, and usefulness.

Week 2 — standardize what works

Save the best prompts and examples. Turn a stable workflow into a reusable skill if useful. Check that another PSA member can retrieve the same source and open the resulting deliverable.

Week 3 — add one useful extension

Choose either a maintained evidence comparison, an interactive pilot-cost explainer, or a navigator demonstration. Use synthetic information for patient-related prototypes. Add a new plugin only if it solves a specific access or workflow need.

Week 4 — decide what to keep

Compare usefulness and rework against observed usage. Keep Astra for work where its added capability matters. Use lighter models for repeatable tasks that meet the same acceptance standard. Check actual billing rather than assuming a token estimate is an invoice.

Your recurring review questions

43. Did the result answer the decision PSA actually needed to make?
44. Were the sources current, accessible, and sufficient?
45. Could colleagues understand and use the output?
46. How much human correction was needed?
47. What should be reused, changed, or discontinued?

29. Plain-language glossary

Term	Plain meaning
Composer	The box where you type or dictate a prompt.
Context	The information available for the current task.
Token	A small unit of information read or written by a model.
Reasoning effort	The amount of model analysis allocated to a task.
Cache	Previously processed information reused to avoid some repeated processing.
Compaction	Condensing earlier working context so a long task can continue.
Agent	A model working with tools toward a task outcome.
Subagent	An additional agent assigned a defined part of the work.
Connector	An authorized route to an external service’s data or actions.
Skill	Reusable instructions and resources for a task.
Sandbox	An environment that limits what tools can access or change.
Rollout	A staged release; accounts may receive a feature at different times.
Overage	Eligible usage beyond the available allocation under the billing arrangement.
BAA	Business Associate Agreement: a contract relevant to covered handling of protected health information.

You do not need to configure APIs, MCP servers, or local policy files to ask Work for PSA reports. Those become relevant if PSA commissions a custom tool or a managed technical deployment.

Official sources: [3] · [13] · [21] · [29]

30. Official sources: model and settings

All sources below were opened or retrieved for this guide. URLs lead to living documentation and may change after September 5, 2026. Recommendations and PSA workflow designs are the guide's own analysis.

[1] GPT-6 Astra model guide

[2] Models and reasoning controls

[3] GPT-6 Astra API specifications

[4] Use ChatGPT: Chat, Work, Codex

[5] Desktop app settings

[6] Personalize ChatGPT

[7] Memories

[8] Projects and chats

[9] Work with files

[10] Pricing and usage

[11] Speed and Fast mode

[12] Workspace model availability

[13] Work usage and cost

[14] Workspace analytics

[15] ChatGPT Work overview

[16] Web search

[17] Browser and cloud browser

[18] Computer History

Scope note: this reference focuses on the end-user and small-team decisions relevant to PSA. Developer mechanics are explained only where they clarify a change or a possible future custom system.

31. Official sources: tools and workflows

[19] ChatGPT Voice

[20] Scheduled tasks

[21] Skills and plugins

[22] Plugins

[23] Sites

[24] Visualizations

[25] Image generation

[26] Record and Replay

[27] Long-running work

[28] Notifications

[29] HIPAA configuration for Codex

[30] What’s new: recent product releases

[31] Async tool calling

[32] Mid-turn steering

[33] Misalignment monitoring

[34] Get started with Work

[35] Prompting

[36] Usage limits and spend controls

[37] Roles and workspace permissions

Edition: 1.0 • September 5, 2026. Recheck model availability, pricing, administrative controls, and interface paths before relying on them after a material product update.