



Product Policy Document: ChatGPT Business (chatbox)

Product and vendor: ChatGPT Business, OpenAI. Until 2025 OpenAI called this tier ChatGPT Team; vendor materials and older firm documents that say "Team" mean this tier.

Way of working: Chatbox. A conversation window in the browser or the desktop app; you type or paste, the model answers.

Status: Ready to adopt, subject to the signup-time confirmations in Section 3. The evidence grade is published, not contract: OpenAI's enterprise privacy page states the commitments, and the binding layer must be confirmed at signup. The binding layer is the Business Terms plus a signed DPA. A DPA, a Data Processing Addendum, is a short attachment to the service contract that states, as enforceable terms, what the vendor may do with the firm's data: process it only on the firm's instructions, keep it confidential, secure it, report breaches, and delete it on request. A privacy page is documentation the vendor can rewrite; the DPA is a signature.

Implements: Model Firm AI Use Policy (42/02). Exhibits: 09 (OpenAI enterprise privacy page, updated 2026-01-08, retrieved 2026-08-07); 11 (Wave 1 exhibits pass memo).

Document date: 2026-08-07.

1. What this is

ChatGPT Business is OpenAI's workspace product for small organizations. The firm buys seats; everyone works inside a shared workspace with central billing and admin controls, separate from anyone's personal ChatGPT account. It is a chatbox: the simplest way of working, and the one most people already know.

Where the data goes: everything typed or pasted into the chatbox leaves the building and is processed on OpenAI's cloud. On this tier, OpenAI commits that business data is not used to train its models by default, and that deleted or unsaved conversations are removed from its systems within 30 days. There is no zero data retention option on this tier; OpenAI offers ZDR only on eligible API endpoints, which is the arrangement a harness uses, not a chatbox (Exhibit 09).

The 30-day deletion commitment carries a qualifier worth reading once in full: conversations are removed "unless longer retention is required by law, or is reasonably necessary to protect our services or any third party from harm." That is the safety-window pattern the policy's Section 3 vetting checklist anticipates: even good terms can hold data longer for legal or trust-and-safety reasons. Know it; do not be surprised by it.

2. The Approved Tool List row (implements Section 3)

ChatGPT (chatbox). Tier / account type: Business workspace, per seat | **Training disabled:** Yes, by default (Exhibit 09) | **Memory off / scoped:** Off in the approved configuration [confirm the workspace memory control at signup; see Section 3, step 4] | **Retention terms:** Admin-controlled; deleted or unsaved conversations removed within 30 days, subject to legal and safety holds (Exhibit 09) | **Permitted input tier:** Tier 2 | **Agentic capabilities:** None in the approved configuration (apps, GPTs, and agent features off or admin-restricted) | **Approved date:** [date of adoption] | **Notes:** No ZDR on this tier. DPA available on request; execute at signup. Formerly named "Team."

3. Required setup (implements Sections 3 and 4)

Signup and configuration, in order. Steps 1 through 7 are the Administrator's; steps 8 and 9 are every user's.

1. **Buy the Business workspace.** Per-seat, self-serve. Do not use consumer accounts (Free, Plus, Pro) for client work under any configuration; they are Tier 3 by definition under Section 4 of the policy.

2. **Execute the DPA.** OpenAI executes a Data Processing Addendum for ChatGPT Business on request through its DPA form (Exhibit 09). This is the step that raises the training and retention commitments from published to contract. Do it before the first client input, and file the executed copy with this document.
3. **Read the Business Terms at signup** and confirm they match the enterprise privacy page on the four load-bearing points: no training by default, 30-day deletion, admin export and delete rights, human review limited to abuse and support cases. Use alone may constitute assent, so read before use, not after (policy Section 3, vetting item 1).
4. **Turn memory off for the workspace.** The policy requires memory features disabled for client work unless the Tool List records a matter-scoped configuration (Section 4). Exhibit 09 does not document the current memory controls on the Business tier, so this is a signup-day confirmation: locate the workspace-level memory and personalization controls, disable them, and record the setting path and date here. If memory cannot be disabled workspace-wide, the row's approval is conditional on per-user verification at training.
5. **Disable public sharing of GPTs and restrict apps.** Workspace members may build internal GPTs; the same data commitments apply inside the workspace (Exhibit 09). External publication of GPTs and third-party app connections are off unless the Administrator approves a specific one; an app that reaches internal sources is a new data path and gets vetted like a tool.
6. **Confirm opt-in data sharing is off.** OpenAI's no-training default has one exception: data explicitly shared through opt-in feedback mechanisms (Exhibit 09). Instruct users not to opt in, and confirm no workspace-level sharing setting is enabled.
7. **Set up access.** Seats for covered personnel only. Use single sign-on (SSO), which lets everyone log in with their firm account instead of a separate password, if the firm has a central login system. Departing personnel lose seats in offboarding, same day.
8. **Users: work signed in to the workspace, never a personal account.** The fastest way to violate this policy is to paste client text into the personal ChatGPT tab that is already open. Log out of personal accounts on work machines, or use separate browser profiles.
9. **Users: complete training** on this document and the Placeholder Card before first client use (policy Section 6).

4. What may go in (implements Section 4)

- **Tier 1** (public information, general legal questions, anonymized hypotheticals): permitted from day one, on any seat.
- **Tier 2** (client documents and matter facts): permitted once the Section 3 steps are complete, the DPA is executed, and no client instruction forbids it. Until then, Tier 1 only.
- **Tier 3** (protective-order material, client-restricted information, statutory confidentiality regimes): prohibited absent informed client consent and Administrator approval, exactly as the policy states. The chatbox is rarely the right venue even with consent.

This product's specific hazard is casual pasting. The chatbox is frictionless: the same window handles a public-law question and, one paste later, a privileged settlement memo. The tier system is enforced at the keyboard, one input at a time, and the Placeholder Card is the working tool: client identifiers, case numbers, and distinctive facts come out or become placeholders before the text goes in, whenever Tier 1 anonymization is the basis for using the tool.

Human eyes exist on the other side. OpenAI's access to stored conversations is limited to authorized employees and specialized third-party contractors, for engineering support, abuse investigation, and legal compliance (Exhibit 09). Automated classifiers run on business data as well. This is a normal enterprise posture,

but it means a flagged prompt can be read by a person who is not at this firm. Write prompts accordingly.

5. Verification (implements Section 5)

Unchanged, by design: the four-step citation ladder, the statute check against official text, and the no-delegation rule apply to chatbox output exactly as to everything else.

Two product notes. First, grounding: the chatbox supports closed-corpus practice only manually, by pasting the statute, the opinion, or the authority file into the conversation and instructing the model to cite only from the provided material; prefer that practice for anything court-bound. Second, documents: the chatbox does not track changes. Output pasted into a client document is a new draft; keep the before version and run a compare, per Section 5's reviewable-record rule.

6. The record (implements Sections 8 and 10)

When a session ends, the conversation exists in one place: OpenAI's servers, under the workspace's retention settings, with the 30-day purge running once a conversation is deleted or unsaved. The firm's record does not live there.

The habit that closes the gap: **export before purge**. For any session that materially assisted work product, the user copies the session (or the operative portion) into the matter file the same day, per Section 10's documentation rule. Workspace admins can view, access, export, and delete end-user conversations (Exhibit 09), which supports supervision under Section 6 and preserves an export path if a user departs; it is not a substitute for the matter-file habit, because admin export is occasional and the matter file is the record.

Prompts and sessions are attorney work product and are stored in the matter file, not in personal accounts (policy Section 4). Billing: per-seat subscription is overhead under Section 8; nothing about this product supports pass-through charges.

7. Vendor facts and evidence grades

No training on business data by default. Source: Exhibit 09 | Grade: Published; contract once DPA executed | Date: Page updated 2026-01-08, retrieved 2026-08-07

Deleted or unsaved conversations removed within 30 days, with legal and safety exceptions. Source: Exhibit 09 | Grade: Published | Date: Same

Admins can view, access, export, delete user conversations. Source: Exhibit 09 | Grade: Published | Date: Same

OpenAI-side human access limited to employees and bound contractors, for support, abuse, legal compliance. Source: Exhibit 09 | Grade: Published | Date: Same

SOC 2 Type 2 audit completed for ChatGPT Business (an independent auditor's report on the vendor's security controls, observed over time). Source: Exhibit 09 | Grade: Published | Date: Same

No ZDR on the chatbox tier (ZDR is API-only, by request, eligible endpoints). Source: Exhibit 09 | Grade: Published | Date: Same

Not established: the text of the Business Terms and the DPA (the contract layer); the current workspace memory controls; the workspace retention-setting granularity. All three are signup-day items in Section 3. **The ask for the vendor, if the firm wants a higher grade:** written confirmation that the enterprise privacy page commitments are contractual for the firm's workspace, and the current documentation for workspace-level memory and retention controls.

8. Re-vetting triggers (implements Section 12)

Any of the following voids this document's conclusions until the Administrator re-checks the cited exhibit against the live page and updates the row:

1. **Tier renames or repackaging.** OpenAI renamed this tier once already (Team to Business). A rename usually accompanies terms changes; re-pull Exhibit 09's source page.
2. **Retention terms change**, or the 30-day window or its exceptions are reworded.
3. **Memory or personalization defaults change** for Business workspaces. This is the likeliest silent change; check at every annual policy review and after any vendor announcement about memory.
4. **Training default changes**, or new opt-in mechanisms appear in the workspace UI.
5. **Apps, GPT store, or agent features become enabled by default** in Business workspaces.
6. **DPA terms change** at renewal.
7. **Acquisition or material corporate change** at OpenAI.