

Product Policy Document: Claude Team (chatbox)

Product and vendor: Claude, Team plan, Anthropic. Anthropic's help center groups its commercial plans (Team and Enterprise) under the name "Claude for Work."

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 position is unusually good for a chatbox: the no-training commitment appears in Anthropic's Commercial Terms of Service as a prohibition on Anthropic itself, and those terms incorporate the DPA automatically.

Implements: Model Firm AI Use Policy (42/02). Exhibits: 06 (Anthropic Commercial Terms, retrieved 2026-08-07); 10 (Claude pricing page, Team and Enterprise sections, retrieved 2026-08-07); 12 (Covered Models retention article, retrieved 2026-08-07); 13 (Anthropic "Using Claude for Legal Work" article, June 2026, retrieved 2026-08-07); 11 (Wave 1 exhibits pass memo).

Document date: 2026-08-07.

1. What this is

The Claude Team plan is Anthropic's workspace product for organizations of 2 to 150 seats. The firm buys standard or premium seats; everyone works in a shared workspace with central billing, single sign-on, and admin controls over connectors. As a chatbox it works exactly like the ChatGPT document (49/03) describes: text typed or pasted into the window leaves the building and is processed on Anthropic's cloud.

Two things distinguish the Claude paper. First, the no-training commitment is written as a contract prohibition, not just a settings default: Anthropic's Commercial Terms state that "Anthropic may not train models on Customer Content from Services" (Exhibit 06, Section B), and the pricing page carries "No model training on your content by default" on the Team plan itself (Exhibit 10). Second, the DPA is incorporated into the Commercial Terms by reference (Exhibit 06, Section C), so there is no separate form to chase at signup, unlike OpenAI, where executing the DPA is a distinct step.

One caution runs the other way: a Team seat is a bundle. It includes Claude Code, Cowork, and other agentic tools alongside the chatbox. **This document approves the chatbox only.** Each agentic tool in the bundle is a different way of working and needs its own Approved Tool List row before anyone uses it for client work.

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

Claude (chatbox). Tier / account type: Team plan, standard or premium seats | **Training disabled:** Yes; contract prohibition (Exhibit 06 Section B) plus no-training default (Exhibit 10) | **Memory off / scoped:** Off in the approved configuration [confirm the memory control scope at signup; see Section 3, step 3] | **Retention terms:** Chats retained until deleted; deletion removes them from the visible history and triggers backend deletion; organization-wide retention windows are listed as Enterprise-only [confirm at signup; see Section 7] | **Permitted input tier:** Tier 2 | **Agentic capabilities:** None in the approved configuration. The bundled Claude Code and Cowork are separate tools needing separate rows | **Approved date:** [date of adoption] | **Notes:** DPA incorporated by reference in the Commercial Terms. Audit logs and the Compliance API are Enterprise-only; see Section 6

3. Required setup (implements Sections 3 and 4)

Steps 1 through 6 are the Administrator's; steps 7 and 8 are every user's.

1. **Buy the Team plan** (2 to 150 seats; standard seats for most personnel, premium for heavy users). Consumer plans (Free, Pro, Max) are prohibited for client work: they are governed by the Consumer Terms, not the Commercial Terms, and consumer training is opt-out rather than off (Exhibit 10, models and usage table).
2. **Read the Commercial Terms at signup and confirm they govern the Team workspace.** The terms state they cover "Anthropic API keys and any other Anthropic offerings that references these Terms" (Exhibit 06); confirm at signup that the Team plan checkout references them. The DPA arrives with them by incorporation. File a copy of both, with the date, alongside this document.
3. **Turn memory off for client work.** Claude now has memory across conversations (Exhibit 10 lists it as a plan feature), which is exactly the cross-matter contamination vector the policy's Section 4 addresses. Locate the workspace or per-user memory controls, disable them, and record the setting path and date here. If memory can only be disabled per user, that check joins onboarding and training. Incognito chats exist as a per-session fallback, but a fallback is not a configuration.
4. **Set the connector controls.** Team admins control remote and local connectors (Exhibit 10). Connectors reach other systems (Microsoft 365, and anything added by MCP); each one is a new data path. Default: off. A connector that touches client data gets vetted like a tool before approval.
5. **Confirm the retention posture and record it.** Anthropic's help center states that Team chats persist for conversation history until deleted, and that custom organization-wide retention windows are an Enterprise feature; Anthropic's legal-work article says "Claude for Work administrators can set an organization-wide retention period" without distinguishing plans (Exhibits 10 and 13; see Section 7). Ask the vendor which is current for Team at signup, and record the answer on the row.
6. **Set up seats and SSO.** Covered personnel only; SSO via the firm's login; seats close at offboarding, same day.
7. **Users: work signed in to the Team workspace, never a personal Claude account.** Personal-account Claude on a work machine invites the same wrong-tab paste as personal ChatGPT; use separate browser profiles or log out.
8. **Users: complete training** on this document and the Placeholder Card before first client use. Training covers the bundle rule: the Claude Code icon in a Team seat is not approved by this row.

4. What may go in (implements Section 4)

The three tiers apply exactly as in the ChatGPT document (49/03, Section 4), and so does the hazard: casual pasting, one input at a time, with the Placeholder Card as the working control.

The Claude-specific addition is the bundle hazard: one login opens the chatbox, and beside it Claude Code, Cowork, and connectors. Those are agentic tools; the policy's agentic access limits (Section 4) and the autonomous-action prohibition (Section 5) apply, and none of them is approved by this document. If the firm wants them, they get their own vetting, their own rows, and in the harness case their own document (compare 49/06 for what a vetted harness looks like).

The covered-model safety window, for completeness: Anthropic retains prompts and outputs to its newest covered models for 30 days for safety work (Exhibit 12). On the Team plan this changes nothing, because the chatbox already runs standard retention; the window matters on ZDR channels (the API and the harness), and Exhibit 12 says so expressly. Do not cite the safety window as a reason to prefer or avoid the Team chatbox; it is not operative here.

5. Verification (implements Section 5)

Unchanged: the four-step citation ladder, the statute check against official text, and the no-delegation rule. Grounding is manual, as on any chatbox: paste the statute, the opinion, or the authority file and instruct the model to cite only from the provided material. Projects can hold reference documents across a matter's sessions; if the firm uses Projects for that, the project is matter-scoped, named for the matter, and its contents are treated as Tier 2 inputs.

The chatbox does not track changes. Output pasted into a client document is a new draft: keep the before version, run a compare.

6. The record (implements Sections 8 and 10)

Same gap, same habit as the ChatGPT document: the conversation lives on the vendor's servers, so **export before purge**, into the matter file, the same day, for any session that materially assisted work product.

One asymmetry to know: audit logs and the Compliance API are Enterprise-only (Exhibit 10). On the Team plan the firm has no vendor-side audit trail of who did what, so the matter-file habit and Section 6 supervision carry the full weight of the record. If the firm's supervision needs outgrow that, the Enterprise plan is the vendor's answer and a new vetting is ours.

Billing: per-seat subscription is overhead under Section 8. Premium seats are still overhead; a seat upgrade is not a client charge.

7. Vendor facts and evidence grades

Anthropic may not train models on Customer Content. Source: Exhibit 06, Section B | **Grade:** Contract (upon assent to Commercial Terms) | **Date:** Retrieved 2026-08-07

No model training on content by default, Team plan. Source: Exhibit 10 (pricing page) | **Grade:** Published | **Date:** Retrieved 2026-08-07

DPA incorporated into Commercial Terms by reference. Source: Exhibit 06, Section C | **Grade:** Contract | **Date:** Retrieved 2026-08-07

Team chats persist until deleted; backend deletion follows; custom retention windows Enterprise-only. Source: Exhibit 10 (help-center facts noted at retrieval) | **Grade:** Published | **Date:** Retrieved 2026-08-07

Covered-model 30-day safety window; inapplicable to standard-retention surfaces like the Team chatbox. Source: Exhibit 12 | **Grade:** Published | **Date:** Retrieved 2026-08-07

Audit logs, Compliance API, custom retention: Enterprise-only. Source: Exhibit 10 | **Grade:** Published | **Date:** Retrieved 2026-08-07

Anthropic's position that its commercial offerings can be configured to maintain privilege, with post-Heppner work-product cases. Source: Exhibit 13 | **Grade:** Vendor commentary; every cited case must be independently verified before any reliance | **Date:** Article dated June 2026

One open discrepancy, deliberately flagged: Exhibit 10's plan grid lists custom retention controls as Enterprise-only, while Exhibit 13 says "Claude for Work administrators can set an organization-wide retention period" without limiting that to Enterprise. The safe reading is that Team retention is not configurable; the signup-day ask is to have the vendor state which is current, in writing. **The ask for the vendor, generally:** written confirmation that the Team workspace runs under the Commercial Terms and DPA, the current memory-control documentation, and the current Team retention posture.

8. Re-vetting triggers (implements Section 12)

1. **Memory defaults change.** Anthropic has been expanding memory features across plans; this is the likeliest silent change. Check at every annual review and after any memory announcement.

2. **The consumer/commercial boundary moves** (which plans the Commercial Terms cover, or a new plan name; compare OpenAI's Team-to-Business rename).
3. **Retention terms change**, or Team gains or loses retention configurability (resolve the Section 7 discrepancy whenever the page changes).
4. **New covered-model designations** or changes to the safety-window terms (Exhibit 12's source page).
5. **Bundled agentic tools become default-on** in Team workspaces.
6. **Connector or MCP defaults change**.
7. **Acquisition or material corporate change** at Anthropic.