

BRIEFING NO. 8

Claude and ChatGPT read your chats.

Not a person, usually. But the switches that decide who does are on by default.



Most people assume a chat with AI works like a phone call: you talk, it answers, it's gone. It isn't. A question typed into ChatGPT or Claude on a regular account makes a series of stops on its way back to you, and most of them are still there after you close the tab.

Ruben Hassid, who writes the *How to AI* newsletter, traced the route in a recent issue. Here is his map, shortened for a busy Monday, plus the settings he says to change today.

Nine stops your prompt makes

- 1. Transmission.** It leaves your computer encrypted and is decrypted on the vendor's servers.
- 2. Recording.** The conversation becomes a database entry. That's true on every subscription tier.
- 3. The model.** The model itself doesn't remember you. The system wrapped around it does.
- 4. Safety scanning.** Automated filters read every prompt. Flagged chats can be kept up to two years; the scores up to seven.
- 5. Human review.** Safety teams and contractors sample conversations. A thumbs-up or thumbs-down sends that exchange to training.
- 6. Storage.** History, memory, and backups persist. In 2025 a court ordered OpenAI to produce 20 million consumer chats for a lawsuit.
- 7. Deletion.** "Delete" schedules removal in about 30 days, unless the data is under legal hold or was already copied into a training set.
- 8. Training.** ChatGPT and Claude both ship with the "improve the model" switch on by default. De-identifying removes your name, not your ideas.
- 9. Side doors.** Shared links are public web pages; Google indexed about 100,000 ChatGPT chats in July 2025. Incognito hides a chat from your sidebar, not from your admin.

Ten minutes of settings changes

Hassid's checklist, in the order that matters most. Menu names move around, so look for the wording rather than the exact path.

- 1. Turn off data training.** ChatGPT: Settings → Data controls. Claude: Settings → Privacy. This is the one switch that matters most.
- 2. Turn off Memory.** What the AI remembers about you is training material unless the toggle is off.
- 3. Delete old share links, and stop making new ones from personal accounts.** Settings → Privacy / Shared links.
- 4. Skip the thumbs-up and thumbs-down.** Rated exchanges are routed to training.
- 5. Use Temporary Chat or Incognito for anything sensitive.** It's in the new-chat menu.
- 6. Remove "ChatGPT sidebar" browser extensions.** Several have leaked user data.
- 7. Move the team to a Business or Team plan (two-seat minimum).** Business data is excluded from training by default.
- 8. Using Codex? Turn off the second switch too.** It covers your whole working environment, not just the chat.
- 9. Keep real secrets off personal accounts entirely.** Client records, contracts, financials, personnel files, source code.

What the settings don't fix

The training toggle closes one door. Stops one through seven are still open: the prompt still leaves your office, still gets recorded, still gets scanned, still sits in someone else's backups and legal holds. For everyday questions that's fine. For a contract, a tax return, or a personnel matter, it's worth pausing before you paste.

A useful test from an earlier briefing: sort the information first. If it's in the "any tool is fine" bucket, go ahead. If it isn't, the question isn't which settings to flip. It's whether the chat should be leaving the building at all.

Where this came from

The nine-stop route and the settings checklist are adapted, with thanks, from "Privacy" by Ruben Hassid, *How to AI* newsletter, September 13, 2026 (ruben.substack.com/p/privacy). Retention periods, the court order, and the indexed-chat count are as reported there. Vendor settings change often; check yours.

Want to learn more about keeping your data private? Email PrivateAI@sterling.net