

AI Privacy Overview

What happens to your data, in plain English · October 2026

→ What happens when you hit send

01

Transmission

Your text is encrypted and sent from your device to the provider's servers.

02

Processing

The model writes a response. A log of the exchange is created.

03

Storage

That log is kept, usually until you delete the chat. Gemini auto-deletes after 18 months.

04

De-identification

Automated filters attempt to strip personal details.

05

Training

What's left may be used to improve future models.



Step 5 is a one-way door. Once your data has trained a model, there is no undo. Not by deleting the conversation, not by closing the account. Everything else on this page is about not reaching step 5 with something you'd regret.

⚙️ Three different things people call "privacy"

Privacy isn't one switch, it's three: **storage**, **human review** and **training**. People collapse them into one word and then assume the worst. You have full control over one, partial control over the second, and very little over the third.

Storage

What it means: a copy of your conversation sits on the provider's servers.

What it does not mean: that anyone reads it, or that it trains anything. It is filing, not reading.

Your control: Very little. It happens on every platform and every tier. You can shorten how long on some, and delete individual chats, but you can't switch it off.

Human review

What it means: in specific situations a person may read a conversation. Not at random, and not routinely. The situations:

- You click thumbs up or thumbs down on a chat
- You contact support and they troubleshoot
- You leave "training on" enabled and it is sampled
- Safety systems flag it as a possible policy violation
- A court or regulator compels it

Your control: Partial. The first three are user-controlled, the last two are system-controlled.

Training

What it means: your conversation **can** be used to improve future models.

What it does not mean: that every conversation is. It means yours is eligible. This is the one that's permanent once it happens.

Your control: Full. A toggle on personal accounts, a contract on business ones.



Turning off training does not mean your data isn't stored, and it doesn't guarantee that no human will ever read it. **A Team, Business or Enterprise account is the only thing that guarantees your data is never used for training** – with one exception, the feedback button, on the next page.

Two things that override every setting you choose

The law / court order

All three providers state plainly that they will retain data, and hand it over, when legally compelled – a subpoena, a warrant, a court order, a regulatory demand.

A court can also order data preserved **past the deletion date the provider advertises**, and this has already happened to a major provider's consumer users. Your retention setting is a company policy, not a legal shield.

The feedback button

Clicking thumbs up or thumbs down sends **the entire conversation**, not just the message you rated.

On personal ChatGPT and Gemini accounts this is an explicit exception to your training opt-out. On Claude, clicking it counts as consent – and it applies **even on Team and Enterprise accounts**, retained for five years. Owners can switch rating off under Organization settings → Data and Privacy.

If a conversation is sensitive, don't rate it.

Turn training off today - every account, every person

Claude	Settings → Privacy → "Help Improve our AI models" → off On by default. You're asked to choose at signup, but the option is pre-selected – clicking through accepts it. Leaving it on extends retention from 30 days to five years .
ChatGPT	Settings → Data Controls → "Improve the model for everyone" → off On by default for Free, Go, Plus and Pro . Paying more does not turn it off for you.
Gemini	Settings → Keep Activity → off, or set auto-delete to 3 months On by default on every consumer tier. Note that history and training are one switch – turning it off also stops saving your chats.

A toggle is not a contract

On a personal account, training is a **setting**. The provider can change the default in a policy update and you have no remedy – one major provider has already revised its consumer terms and made every existing user re-choose. On a business account it is a **contractual commitment** with a Data Processing Addendum behind it, and it cannot change unless both parties agree. Upgrading isn't buying features. It's buying legal protection over your data.

-  **The one rule worth memorizing.** Anything you wouldn't email outside the company goes on the business account – customer data, financials, source code, legal documents, HR records. Personal accounts are for public and low-stakes work. And passwords, API keys and account numbers go into no tier, ever.

What deleting a chat actually does

Claude	Purged within 30 days and not used for future training. Rated chats are kept five years.
ChatGPT	Removed within 30 days unless a legal hold applies. Temporary chats are still kept up to 30 days.
Gemini	Removed from your account, but human-reviewed chats are kept up to three years.

Consumer accounts vs business accounts

	CLAUDE	CHATGPT	GEMINI
Training, consumer tiers	On by default, all tiers	On by default, all tiers	On by default, all tiers
Consumer retention	Until deleted, then purged within 30 days; 5 years if training is on	Until deleted, then purged within 30 days	Auto-deleted after 18 months by default
Human review possible	Yes – flagged content and feedback	Yes – abuse, support, legal, feedback	Yes – sampled chats kept 3 years
Training, business tiers	Never, by contract	Never, by contract	Never, by contract
Data Processing Addendum	Team and Enterprise	Business and Enterprise	Workspace Business and up
HIPAA agreement (BAA)	Enterprise only	Specific products only	Any plan, with admin-accepted BAA

All three default to training on. Claude asks at signup rather than switching it on silently – but the option is pre-selected.

Choosing a business plan

Do you need to set how long conversations are kept? **Enterprise**

Do you need a HIPAA agreement? **Enterprise, specific products**

Do you need audit logs or a compliance export? **Enterprise**

None of the above? **Business / Team is enough**

lumioconsulting.com/ai-privacy

The full guide, with plan-by-plan tables for Claude (Team vs Enterprise) and ChatGPT (Business vs Enterprise), fact-checked and dated.

The official plan pages

claude.com/pricing · chatgpt.com/pricing · Google Workspace Generative AI Privacy Hub. Re-check before you sign; tiers change often.

Before you point an agent at your files

Local files are not local

When an AI agent works on files on your computer, the contents are sent to the provider's servers to be processed. Coding tools send only the files needed as context – but that still leaves your machine.

"It's on my laptop" is not a privacy control. The account tier decides what happens next, same as chat.

Prompt injection

Only a risk once an agent can read your email, documents or the web. Hidden instructions buried in a document or web page can hijack the agent and make it leak data.

It is the **number one risk** on the industry's list for AI applications – and giving an agent broader access than the task needs sits at number six. **Grant the narrowest access that does the job.**