AI · CONFIDENTIALITY

Before You Paste That Intake Into *ChatGPT*

What every attorney should know about public AI versus enterprise-grade AI, and where your client's data actually goes the second you hit enter.

By **Renee Waite**, Founder, Simplarity

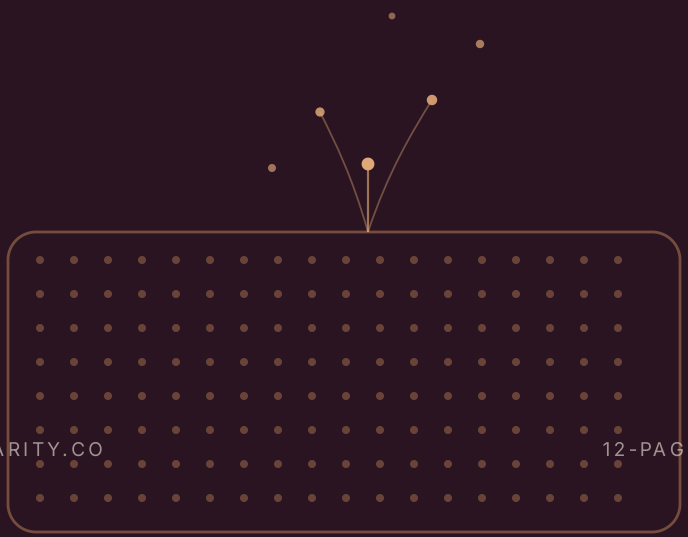

THE REAL QUESTION

The question is not which AI is best.

It is where your client's information goes the moment you hit enter.

Most firms never ask. Someone opens a consumer chatbot, pastes an intake, a draft declaration, a settlement figure, a tax transcript, and moves on. The work gets faster. The exposure stays invisible.

Here is the part nobody mentions at the CLE. On most consumer AI plans, including the ones you pay for, your inputs can be stored for years and used to train the next version of the model. That is not a worst-case reading of the fine print. After the 2025 policy changes across the major tools, it is the default setting.

The fix is not to ban AI. It is to understand the one distinction that decides whether a tool is safe to point at a client matter.

> **Paying twenty dollars a month does not buy you privacy.** It buys you features. On most paid individual plans, your data still trains the model unless you have personally switched it off.

THE FRAMEWORK

Three questions every tool has to answer before it touches a client matter.

01 **Does it train on my inputs?**
On most consumer tiers the answer is yes by default, and opting out is on you. On a true business or enterprise contract, no training is the default and it is written into the terms.

02 **How long is my data kept, and can I get Zero Data Retention?**
ZDR means your prompt is deleted right after it is processed. Never written to logs. Never seen by a human reviewer. It is the gold standard, and it is almost always an enterprise or API feature, not a consumer one.

03 **Is there a signed DPA, where does the data live, and do I need a BAA?**
A Data Processing Addendum and U.S. data residency protect everyday client confidentiality, and they apply to every matter you handle, so confirm those first. A Business Associate Agreement is narrower. It governs protected health information, which makes it matter for attorneys whose cases involve medical records. Add the BAA when a matter actually touches PHI.

> If a tool cannot answer all three **in writing**, it does not belong anywhere near privileged material. Not for a draft. Not for a quick cleanup. Not once.

PUBLIC VERSUS ENTERPRISE

The one line that actually matters.

PUBLIC / CONSUMER

- **Built to improve the model.** Your inputs are the fuel.
- Trains on your data by default on most plans. Opt-out is your job.
- Long or open-ended retention. Five-year windows are now common.
- No BAA. Not HIPAA-eligible.
- No admin controls, no audit log, no signed DPA.

ENTERPRISE-GRADE

- **Built to protect the buyer.** No training on your data by default.
- Configurable retention, with Zero Data Retention available.
- BAA available for qualifying use.
- Admin controls, audit logs, a signed Data Processing Addendum.
- You are the customer, not the training set.

Notice what the line is *not*. It is not free versus paid. The privacy boundary sits between consumer and business, and a paid individual plan lands on the wrong side of it.

> **The Pro and Plus trap.** A twenty-dollar individual seat is still a consumer plan. It usually upgrades your speed and your model, not your confidentiality terms.

04

Claude Anthropic

CONSUMER · FREE, PRO, MAX

Since the August 2025 policy change, a single setting, **Help improve Claude**, decides whether your new chats and coding sessions train future models. Left on, training-eligible data can be retained for up to five years. You opt out by hand. Consumer accounts are not HIPAA-eligible and carry no BAA.

ENTERPRISE-GRADE · TEAM, ENTERPRISE, API

Under Commercial Terms, your inputs are never used for training by default. The API deletes inputs and outputs after seven days, extendable to thirty through a DPA. A BAA is available for qualifying healthcare use, and Zero Data Retention is available for qualifying enterprise and API customers.

ONE CATCH WORTH KNOWING

Under Anthropic's BAA, some features are switched off, including web search. The compliant configuration is a narrower tool than the consumer app. Plan your workflow around that, not against it.

Practitioner note. That toggle is easy to miss, and it lives on each individual seat. If your firm runs on Pro, check Help improve Claude on every login today, not next quarter.

05

ChatGPT OpenAI

CONSUMER · FREE, GO, PLUS, PRO

Trains on your conversations by default. You opt out under Settings, then Data Controls. None of these tiers are HIPAA-eligible, and none carry a BAA.

BUSINESS · ENTERPRISE · API

No training on workspace data by default, with admin controls, custom retention, and Zero Data Retention on qualifying terms. A BAA is available on sales-managed Enterprise, Edu, the dedicated ChatGPT for Healthcare plan, and through the API. Note that ChatGPT Business on its own is not HIPAA-eligible.

THE RECEIPT

From May 2025 to October 2025, a federal court preservation order in the *New York Times* litigation required OpenAI to retain Free, Plus, Pro, Team, and non-ZDR API conversations, including deleted chats and Temporary Chats. Enterprise, Business, and Edu workspaces were carved out. On consumer tiers, for those months, **deleted did not mean deleted.**

06

Codex OpenAI

Codex is not a separate privacy regime. It is a coding agent that runs inside whatever plan or contract you sign into, and it inherits those data terms.

RUN UNDER BUSINESS, ENTERPRISE, OR API

No training on your data by default. Same protections as the business plan it sits under.

RUN UNDER PERSONAL FREE OR PLUS

Consumer rules apply. Your inputs can train models unless you have already opted out on that account.

WHY A LAW FIRM SHOULD CARE

Nobody thinks of a coding tool as a place client data goes. Then someone pastes a contract, a cap table, or a client's financials into Codex to have it cleaned up, on a personal plan, and it carries the same exposure as the chatbot with none of the awareness. The tool is only as safe as the account it is signed into.

07

Perplexity

CONSUMER · FREE, PRO, MAX

Training is on by default. You opt out with the AI data retention toggle in Account Settings, and the opt-out only applies going forward. No admin controls or DPA on these tiers.

ENTERPRISE PRO · ENTERPRISE MAX

No training on your data, contractually. Zero Data Retention offered. SOC 2 Type II, a DPA, SSO, admin controls, custom retention. A HIPAA BAA is available on Enterprise tiers only, and as of February 2026, PHI is forbidden without one in place.

THE SOPHISTICATED CATCH

Perplexity routes through other companies' models, including OpenAI, Anthropic, and Google. On Enterprise those come with zero-retention agreements. But its newer multi-model routing API can place you under all of those providers' terms at once. With any tool that routes, you inherit every downstream vendor's policy. **Read the path, not just the label.**

AT A GLANCE

Four tools, one honest grid.

TOOL	CONSUMER DEFAULT: TRAINS ON YOU?	ENTERPRISE: NO TRAINING?	ZDR	BAA
Claude	**Yes**, unless opted out	**Yes**, Commercial Terms	**Yes**, Ent & API	**Yes**, qualifying
ChatGPT	**Yes**, unless opted out	**Yes**, Business, Ent, API	**Yes**, qualifying	**Yes**, Ent, Edu, Health, API
Codex	**Inherits the plan**	**Yes**, Business, Ent, API	Per the plan	Per the plan
Perplexity	**Yes**, unless opted out	**Yes**, Enterprise Pro, Max	**Yes**, Ent & Sonar API	**Yes**, Enterprise only

Current as of May 2026. Vendor terms change often, and they changed materially in 2025. Confirm each tool's live terms before you rely on this.

> The pattern is the same everywhere. **Consumer trains on you. Enterprise does not.** The work is in knowing which seat each person at your firm is actually using.

THE DUTY

This is not a tech preference. It is a duty.

ABA Formal Opinion 512, issued July 2024, did not ban AI. It put generative AI inside the Model Rules and told lawyers what they have to do before they use it. Two duties carry the weight.

Confidentiality · Model Rule 1.6

Before you enter client information into a GenAI tool, you have to evaluate the risk that it could be disclosed to or accessed by others, including whether your inputs are used to train the model. That duty runs to all confidential client information, not just medical records, so the protections that matter most are a no-training commitment and a signed DPA. A BAA is the added layer only when a matter involves protected health information.

Competence · Model Rule 1.1

You are expected to understand, at a working level, how the tool handles data. Not the math. The data path.

"I think it's private because I pay for it" is not a defense.

THE CHECKLIST

Sixty seconds before you paste.

- **Consumer or business?** If it is a consumer plan, assume it can train on you.
- **Is training actually off** on this exact seat, not just the firm's main one?
- **Is a BAA in place** if there is any protected health information in the matter?
- **Do I need Zero Data Retention** for this client, and does this tool offer it?
- **Where does the data physically live,** and does my engagement letter or the client allow it?
- **If the tool routes to other models,** whose terms am I also accepting?

> If you cannot answer these in under a minute, **the tool is using you, not the other way around.** Save this page. Print it. Tape it to the monitor of whoever does intake.

THE WORK BEHIND THE GUIDE

Built by someone who files *the cases.*

Simplarity is AI workflow consulting for immigration practices, led by Renee Waite. The reason this guide names settings and retention windows instead of buzzwords is simple. The work it protects is work that gets done at 2 a.m., not theorized on a panel.

If your firm is using AI and nobody has mapped where the client data goes, that is the place to start.

Book a Simplarity efficiency audit

One audit. Every AI tool your firm touches, mapped against the three questions in this guide.

Systems that simplify. Results that last.