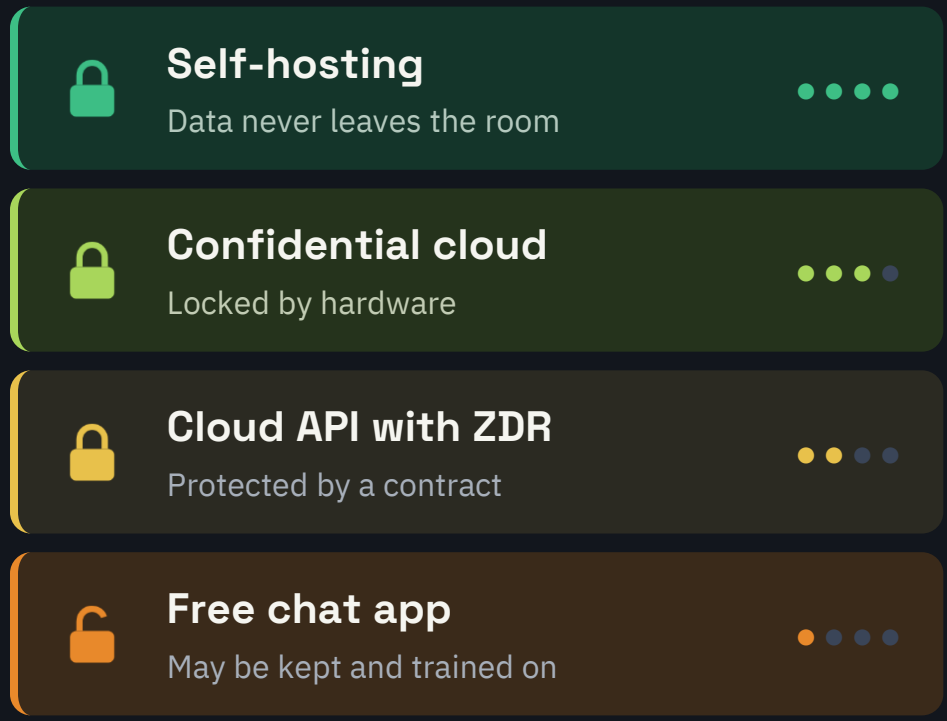
A BEGINNER'S GUIDE

Keeping your data private with AI

Comparing privacy options, data retention policies and provider safeguards



Ibrahim Jaafar

WHY IT MATTERS

Every prompt is data you hand over

1

It can be stored

Copies sit on the provider’s servers for days, months or years.

2

It can be trained on

Your words can help shape the next version of the model.

3

It can be seen

By human reviewers, by attackers after a leak, or through a court order.

What people paste in: source code · client contracts · health questions · passwords · company plans

WHY IT MATTERS

What is at stake

Compliance and the law
GDPR fines reach €20M or 4% of global turnover; the EU AI Act up to 7%.

Personal information
Names, IDs, health and finances. Once leaked, they can’t be taken back.

Company data
Source code, strategy and pricing are your competitive edge.

Client trust and contracts
NDAs often forbid sharing client data with any third party, AI included.

Government and national security
Classified or public-sector data may only go to approved systems.

REAL EXAMPLES

Apr 2023 · Samsung

Engineers pasted internal source code into ChatGPT. Samsung then banned generative AI at work.

Dec 2024 · OpenAI

Italy’s regulator fined OpenAI €15M for training on personal data without a legal basis.

Jan 2025 · DeepSeek

An open database exposed over a million log lines, including user chat history.

ROADMAP

What we'll cover

1	**Where your prompt goes**	Storage, training and who can see it
2	**Privacy options, ranked**	From your own computer to free apps
3	**Zero Data Retention**	What ZDR is, and its catch
4	**Providers and applying**	Who offers what, and how to ask for it
5	**Choosing your setup**	Match the tool to the data, plus daily habits

GLOSSARY

Nine words you'll hear

Prompt
Anything you send: text, files, images

Retention
How long the provider keeps a copy

Training
Using your data to teach future models

Abuse monitoring
Safety checks that keep prompts to catch misuse

ZDR
Zero Data Retention: nothing kept after the reply

Self-hosting
Running the model on hardware you control

Open-weight model
A model you can download and run yourself

DPA
Data Processing Agreement: the legal privacy contract

TEE
Trusted Execution Environment: a hardware-locked box

1

PART 1 OF 5

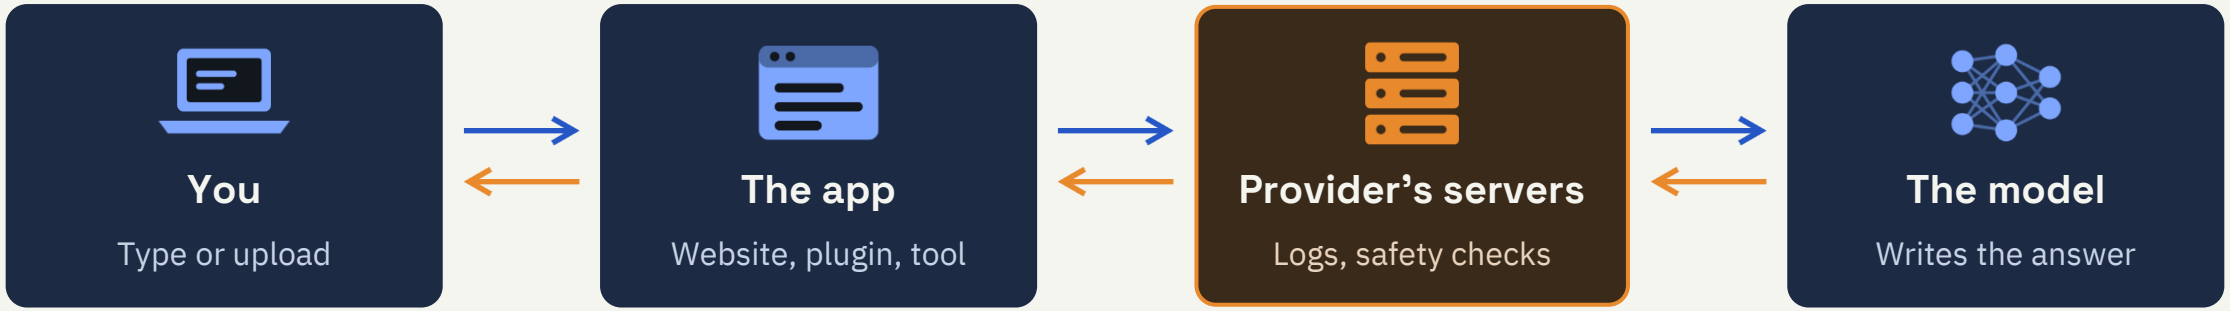
The journey

Where your prompt goes after you press Enter

Storage, training, and who can see it.

1 · THE JOURNEY

From your keyboard to their servers



At every stop, ask four questions:

- Is it stored?
- For how long?
- Is it trained on?
- Who can see it?

THE MOST IMPORTANT IDEA TODAY

“We don’t train on your data” does not mean “we don’t store your data”

Most paid AI services promise no training, but still keep your prompts for up to 30 days to check for abuse.

Cloud AI: same model, different rules per plan

How you use it	Trained on?	How long it's kept	Humans may read?
Free app or free API tier	Often yes	Months to years	Yes, can be reviewed
Paid personal app (Plus, Pro, Max)	Often on by default	Depends on your setting	If flagged
Business plan (Team, Enterprise)	No	Until deleted (admin sets)	If flagged
Standard paid API	No	Up to ~30 days	If flagged
API with ZDR	No	Not kept after the reply*	Only flagged content

Typical pattern across OpenAI, Anthropic, Google and Microsoft. Exact rules differ by provider and change often; see Part 4 for each one. * With exceptions, covered in Part 3.

2

PART 2 OF 5

The ranking

Privacy options, from most to least private

And what each one costs you in quality, money and effort.

2 · THE RANKING

The privacy ladder

#	Option	Privacy	Who can see your data
1	**Your own computer, offline**		Only you
2	**Your own server, in your office**		You and your team
3	**Rented GPU, confidential mode (TEE)**		You; the hardware locks out the host
4	**Rented GPU or private cloud you run**		You, plus the cloud host in theory
5	**Cloud API with ZDR**		Provider, only while answering
6	**Paid API or business plan**		Provider, for ~30 days; no training
7	**Personal app, training turned off**		Provider, by its consumer policy
8	**Free app or free API tier**		Provider, reviewers, future models

What each option costs you

Option	Privacy	Model quality	Cost	Effort
Self-host on your machine	Highest	Good, limited by your GPU	Hardware up front	High
Self-host on a rented GPU	Very high	Very good, big open models	Per GPU hour	High
Confidential computing	Very high, provable	Good, mostly open models	Premium	Medium
Cloud API with ZDR	High, by contract	Best (top cloud models)	Per use + business deal	Medium (approval)
Paid API or business plan	Medium	Best	Per use or per seat	Low
Free apps and tiers	Low	Good, with limits	Free (you pay with data)	None

The trade-off: more privacy usually means weaker models, more money, or more work.

2 · THE RANKING

Self-hosting: most private, not free

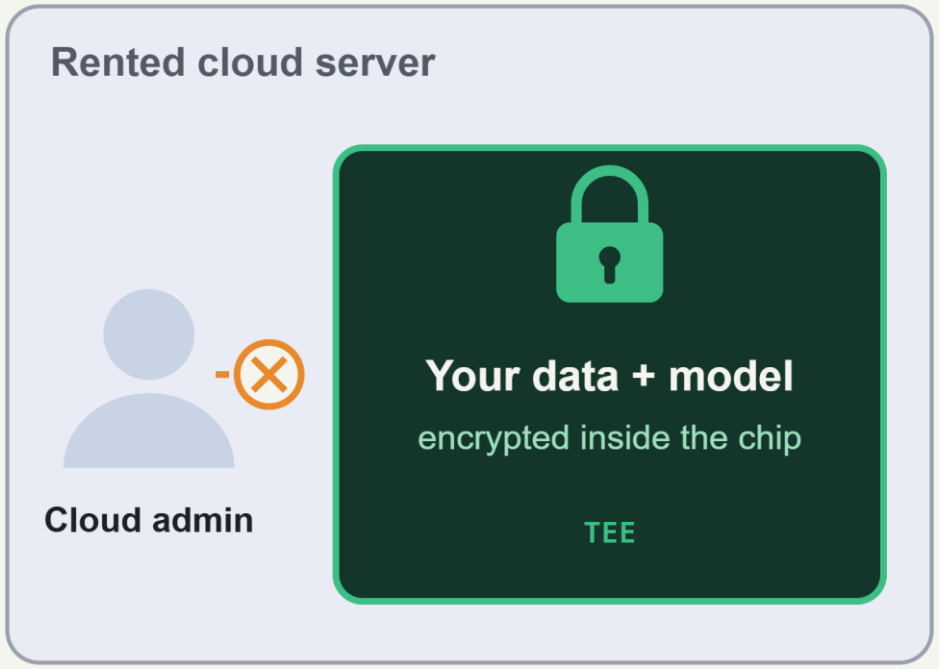
Why it wins

- Data never leaves your machine
- No retention policy to trust
- Works fully offline
- No terms that can change overnight

The catch

- You are now the security team
- Open models trail the best cloud models
- Needs a GPU with enough VRAM
- Setup and updates are on you

Confidential computing: a locked box in the chip



TEE = Trusted Execution Environment
The hardware that makes confidential computing work.

1. The chip **encrypts memory** while the model is running.
2. It signs an **attestation**: proof the box is genuine and untouched.
3. Even the **cloud company can't look inside**.

ZDR: you trust a contract.
Confidential computing: you verify with hardware.

Example: Apple Private Cloud Compute. The catch: mostly open models, at a premium price.

2 · THE RANKING

Six safeguards that work with any AI setup

Send less

Share only the part the AI needs, not the whole file.

Redact

Swap names, IDs and numbers for placeholders like [CLIENT].

Route by sensitivity

Secrets go to a local model; everything else to the cloud.

Guard your API keys

Keep them in a vault, never in code. Rotate them often.

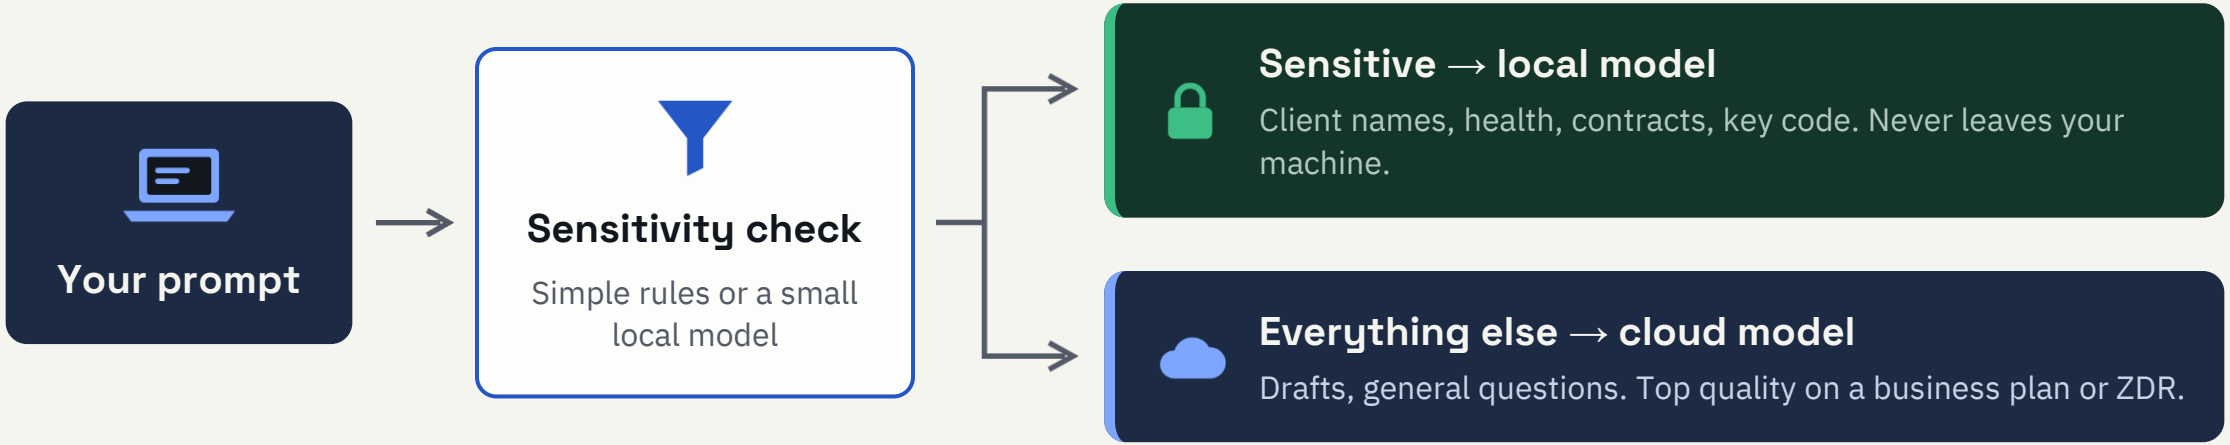
Mind your own logs

Your app may save prompts too. Encrypt them and auto-delete.

Watch the add-ons

Web search, uploads and plugins send data to more places.

A hybrid setup: route by sensitivity



Start simple: “Anything with a client’s name goes to the local model” is already a working rule.

3

PART 3 OF 5

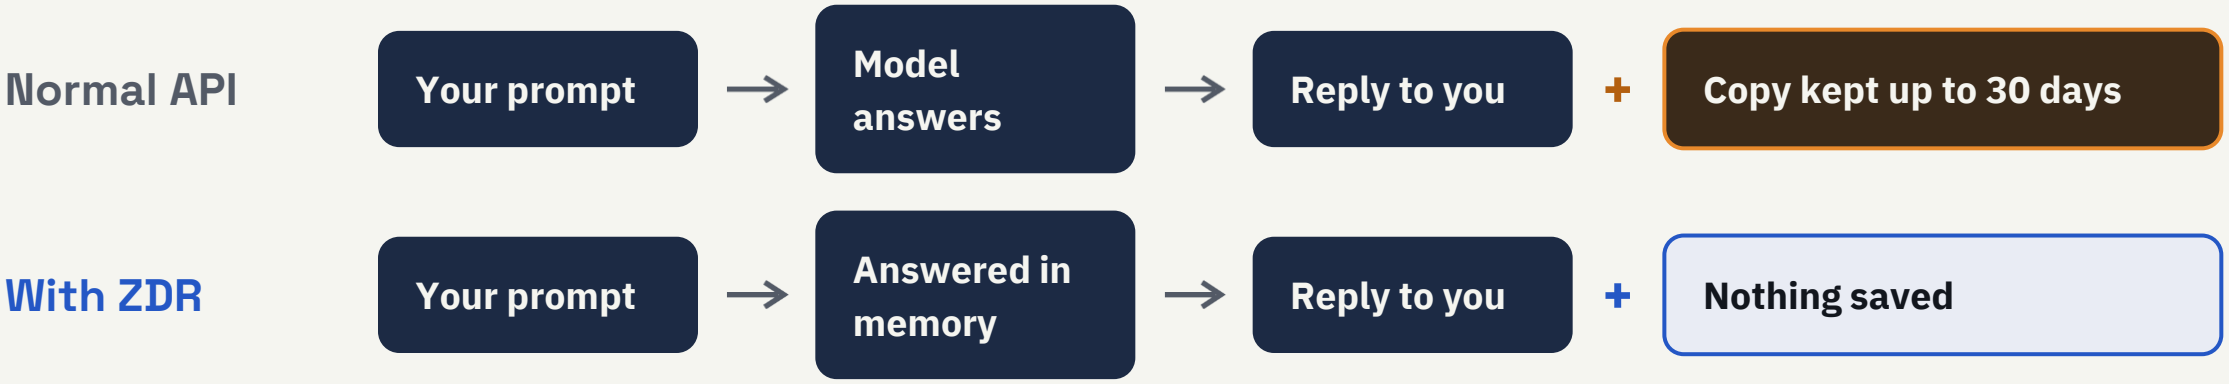
Zero Data Retention

The strongest promise a cloud provider makes

What ZDR is, what it isn’t, and its catch.

3 · ZERO DATA RETENTION

What ZDR changes



What it is

A signed agreement: the provider won’t store your prompts or answers after replying.

What it isn’t

A technology you can check. You trust the contract, not the hardware.

3 · THE CATCH

Zero isn’t always zero

You must apply

Not a switch. It needs approval, usually an enterprise deal through sales.

Only some features

File uploads, batch jobs, caching and saved responses can still store data.

Flagged content stays

Safety flags still apply. Anthropic keeps flagged chats up to 2 years.

Chat apps aren’t covered

ZDR covers the API. ChatGPT or Claude.ai follow their own rules.

Terms can change

A new model can be left out. It happened in 2026 (next slide).

You can’t verify it

Your data is still readable on their servers while it’s processed.

ZDR in the real world

May 2025

Court orders OpenAI to keep all logs

Part of the NYT lawsuit. ZDR customers weren't affected: nothing was stored.

Sep 2025

The order ends

OpenAI says the duty to keep everything ended on 26 September 2025.

Jun 2026

Claude Fable 5 skips ZDR

Launched with required 30-day retention. Existing ZDR deals didn't apply.

Sep 2026

ZDR returns with Fable 5.1

After enterprise pushback, eligible customers get ZDR again.

The lesson: ZDR shields you from leaks and court orders. But it's a contract: re-check it every time you switch models.

4

PART 4 OF 5

Providers

Who offers what, and how to apply

OpenAI, Anthropic, Google, Microsoft and Amazon.

The big providers at a glance

Provider	Trains on paid data?	Kept by default	How to get ZDR
OpenAI API	No	Up to 30 days	Apply via sales, then turn on in organization data controls
Anthropic (Claude API)	No	Up to 30 days	Ask sales or your account rep; enabled per organization
Google Gemini API	Paid: no. Free: yes	Free tier may be human-reviewed	Use the paid tier; for ZDR, use Vertex AI
Google Vertex AI	No	Cache up to 24 h; abuse logs	Turn off caching, request abuse-monitoring exception
Azure OpenAI	No	Up to 30 days (abuse checks)	Apply for “Modified Abuse Monitoring” via your Microsoft rep
Amazon Bedrock	No; model makers can’t see it	Not logged by default	Set retention mode to “none”; approved per model

Policies change often. Checked October 2026; always confirm on the provider’s own privacy page.

How to apply for ZDR, step by step

1 Get a business API account
ZDR is never offered on free or personal plans.

2 List what you use
Models and features: files, batch, caching, web search.

3 Contact sales with a reason
Explain the data: health, legal, finance, client secrets.

4 Sign the paperwork
ZDR amendment and a DPA; a BAA too for US health data.

5 Switch it on, switch storage off
Enable ZDR for the org; turn off caching and saved responses.

6 Verify, document, repeat
Test it, write it down, and re-check for every new model.

Small team? ZDR can be hard to get. Bedrock, Vertex AI or Azure may be easier routes.

4 · PROVIDERS

Eight questions to ask any provider

- [x] Do you train on my data? Can I opt out?
- [x] How long do you keep prompts and answers?
- [x] Can your staff read my data? When?
- [x] Where is it processed? (which region)
- [x] Who else touches it? (subprocessors)
- [x] Which certificates? (SOC 2, ISO 27001)
- [x] Will you sign a DPA (GDPR) or BAA (HIPAA)?
- [x] What happens if a court asks for my data?

4 · PROVIDERS

Laws, and where your data lives

GDPR · European Union

Covers personal data of people in the EU, even if your company is elsewhere. Needs a legal basis and a DPA with your AI provider.

HIPAA · United States

Covers patient health data. Use only “HIPAA-eligible” AI services, with a signed BAA.

Law 81/2018 · Lebanon

Covers personal data handled electronically: clear purpose, informing people, keeping it secure. Narrower than GDPR.

Data residency · everywhere

Where data is processed matters. Many providers let you pick a region, like the EU or US, to match your rules.

Not legal advice. For regulated data, check with a lawyer or your compliance team.

5

PART 5 OF 5

Choosing

Match the setup to the data

Plus the everyday habits that matter most.

How sensitive is the data?

Your data	Good-enough setup
Public Blog drafts, public docs, learning	Any tool, free apps are fine
Internal Meeting notes, plans, everyday code	Business plan or paid API (no training)
Confidential Client files, personal data, contracts	API with ZDR, or a cloud platform in your region
Secret or regulated Health, legal cases, key source code	Self-host, or confidential computing

5 · CHOOSING

Six habits to start today

01

Never paste secrets

No passwords, API keys or card numbers. Ever.

02

Turn training off

Find “improve the model” in your app’s privacy settings.

03

Work account for work

Personal plans put company data under consumer rules.

04

Use placeholders

Write [CLIENT] and [AMOUNT] instead of the real thing.

05

Use temporary chats

For one-off questions that don’t need history.

06

Clean up old chats

Delete what you no longer need. Less stored, less risk.

An AI privacy policy for your team

1 Approved tools only
A short list of AI tools staff may use for work.

2 Data levels, mapped to tools
Public, internal, confidential, secret: where each may go.

3 Company accounts only
No personal plans for work data, ever.

4 Settings set by an admin
Training off and retention limits, on every account.

5 A “never paste” list
Passwords, keys, ID numbers, health and client secrets.

6 One owner, regular reviews
Someone to ask, and a check whenever providers change terms.

Keep it short: a one-page policy people actually read beats a long one nobody does.

RECAP

Five things to remember

1. “No training” is not the same as “not stored”.
2. Self-hosting is the most private. Free apps are the least.
3. ZDR is strong, but it’s a contract with exceptions.
4. Check the terms for the exact plan and model you use.
5. Match the setup to how sensitive the data is.

SOURCES

Read more

OpenAI: response to NYT data demands
openai.com/index/response-to-nyt-data-demands

Anthropic: what ZDR applies to
support.anthropic.com › zero data retention

Google: Vertex AI and zero data retention
docs.cloud.google.com › data-governance

AWS: Amazon Bedrock data retention
docs.aws.amazon.com › bedrock › data-retention

Anthropic: API and data retention
platform.claude.com › api-and-data-retention

Anthropic: how long data is stored
privacy.anthropic.com › how long

Microsoft: Azure OpenAI abuse monitoring
learn.microsoft.com › abuse-monitoring

CX Today: ZDR returns for Fable 5.1
cxtoday.com › anthropic zero data retention