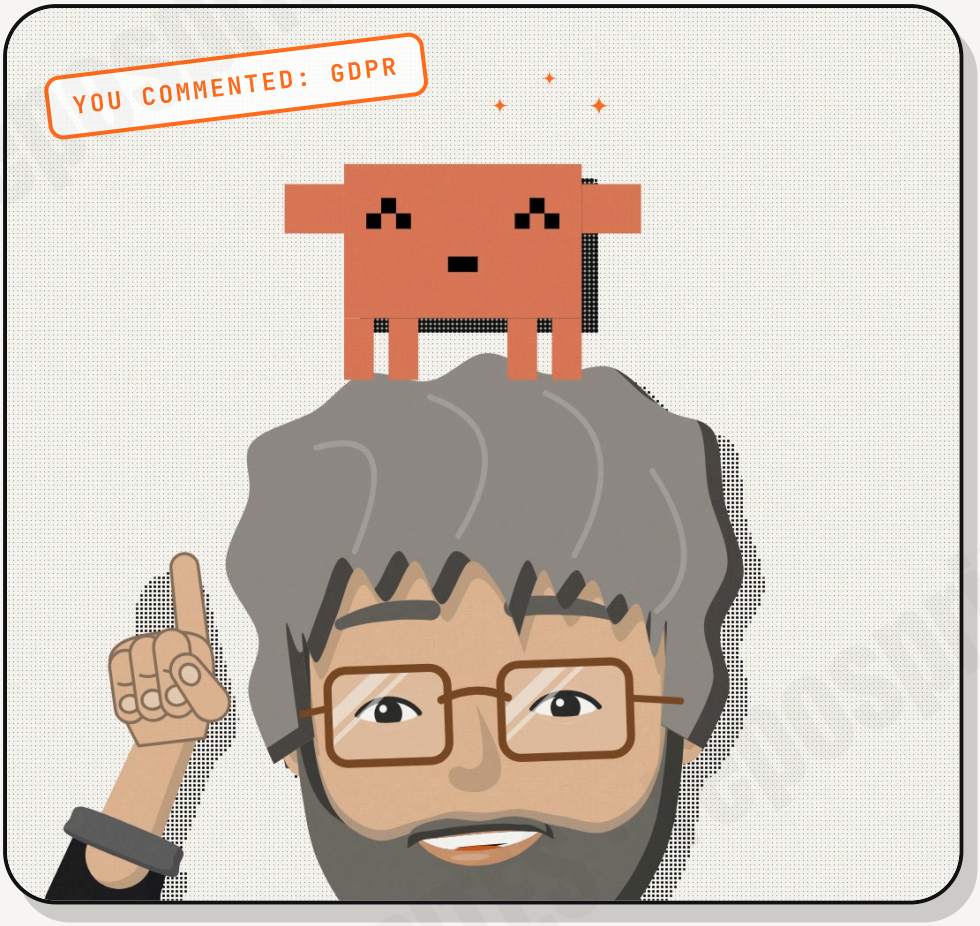


AI and GDPR: tricky, *but solvable*

Part 2 of the agent office. The simple way with Langdock, your own loop with Presidio, ready-made building blocks, and the questions you should ask every provider.

@reposprites

reposprites.com

Thanks for following ✓ · You commented: GDPR ✓

© 2026 REPOSPRITES · @reposprites · reposprites.com · Personal copy, no sharing or resale · ID RS-20261103-ADA04261



THE MENTOR SAYS

Orientation, not legal advice. For sensitive data (health, clients under professional secrecy, employees) talk to your data protection officer.



THE MENTOR SAYS

Independent: I have no connection to Langdock. Everything here is based on my own experience and research.

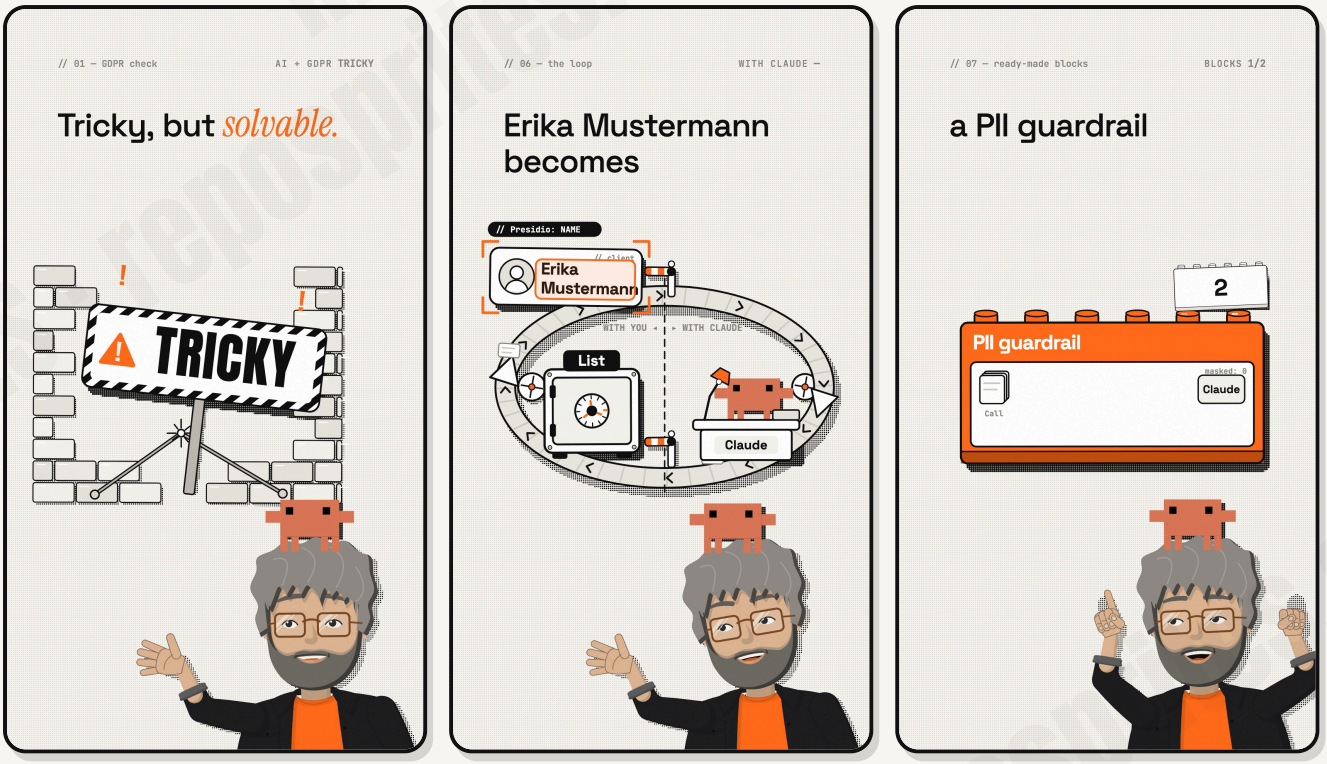
// 01

At a glance

	LANGDOCK	YOUR OWN LOOP
For whom	Teams who want Claude without any tech	Anyone who wants full control over names and client data
Tech needed	None, runs in the browser	Yes: Presidio, a gateway or hooks
What Claude sees	Your input, as you write it	Only placeholders like `<PERSON_1>`
Contract	Data processing agreement built into the terms	Agreement with your model provider is still needed
Cost	Business €25 net per user per month according to the pricing page	Software free (open source), plus model costs and your time
Limit	Backend per Claude model not public, no built-in name masking found	Doesn't catch everything, samples needed

© 2026 REPOSPRITES · @reposprites · reposprites.com · Personal copy, no sharing or resale · ID RS-20261103-ADA04261

// FROM THE VIDEO



// 02

Way 1: Langdock, the simple way

Langdock is an AI workspace from Berlin in the browser: chat, agents, workflows, with Claude and other models. I have no connection to Langdock, this is my own experience and research.

- **Hosting:** Microsoft Azure with servers in the EU. According to Langdock, the app and "most" models run fully in the EU.
- **Data processing agreement:** part of the terms, applies automatically.
- **Professional secrecy:** the Trust Center lists an add-on agreement for lawyers, doctors and tax advisers (German § 203 StGB).
- **Security:** ISO 27001, SOC 2 Type II, AES-256, admin roles, a configurable retention period for chats. According to Langdock, customer data is never used to train AI models. What these badges mean is explained below.
- **MCP:** remote servers only (HTTP or SSE), no local ones.
- **Price:** Business €25 net per user per month according to the pricing page, Business Max €99. A review site saw different numbers, check the price before you buy.

What's open: on the models page every Claude model has an EU icon, but no backend is named. According to the API docs, requests go to AWS Bedrock, Google Vertex or Anthropic's own servers, depending on the model. I found no built-in feature that masks names before sending.

© 2026 REPOSPRITES · @reposprites · reposprites.com · Personal copy, no sharing or resale · ID RS-20261103-ADA04261

The security badges in plain words

Langdock's security page lists four points under audits and certifications. They are good signs. Here is what each one means, and where your own part begins.

BADGE	WHAT IT MEANS	WHERE TO LOOK CLOSER
ISO 27001, "certified"	An outside auditor checked that Langdock runs a working system for information security: risks, access, backups, incidents. Checked again every year.	It only covers what is in the certificate's "scope". The page doesn't name the scope, the auditor or the date.
SOC 2 Type II, "independently audited"	A US audit (AICPA standard). "Type II" means the auditor checked over several months that the controls actually worked, not just that they exist on paper.	The report itself is not public. You usually get it on request under an NDA. It is a security audit, not a privacy or GDPR check.
GDPR, "fully aligned"	Langdock lists GDPR under its audits and certifications and says it is fully aligned: an agreement in the terms, EU hosting, rights for the people whose data it is.	It's their own assessment: there is no official GDPR certificate to show, for Langdock or any other provider. And no tool makes *you* compliant: your legal basis, informing people and what you put in stays your job.
"No model training"	According to Langdock, your data is never used to train or improve AI models. Anthropic and AWS Bedrock say the same for business use.	Not training is not the same as not storing. The model provider can still keep requests for a short time, for example to check for abuse. The page doesn't say how this works for each model.

In short: ISO 27001 and SOC 2 Type II tell you Langdock takes security seriously and lets outsiders check it. Langdock does its part for the GDPR, your own use is still up to you. Get the promises in writing, in the agreement, not only on a web page.

© 2026 REPOSPRITES · @reposprites · reposprites.com · Personal copy, no sharing or resale · ID RS-20261103-ADA04261

Ask Langdock these questions in writing

- [ ] Which backend and region does the Claude model we use run on?
- [ ] Are prompts stored there, for how long, and who can review them?
- [ ] Which sub-processors are involved, and where are they?
- [ ] Is there masking of names or client data before it reaches the model?
- [ ] What exactly does the add-on agreement for professional secrecy cover?
- [ ] Can we see the ISO 27001 certificate (scope, auditor, valid until) and the SOC 2 Type II report?
- [ ] Is "no training" written into our agreement, and does it also apply at the model provider for our model?

// 03

Way 2: your own loop

The idea: names and client data never leave your building in the first place.

1. Your text first goes through a tool on your own computer, for example Presidio.
2. It replaces names, IBANs and ID numbers with placeholders: Erika Mustermann becomes `<PERSON_1>`.
3. The mapping list (`<PERSON_1>` = Erika Mustermann) stays with you.
4. Only the placeholder text goes to Claude.
5. The answer comes back with placeholders, and your loop puts the real names back in.

© 2026 REPOSPRITES · @reposprites · reposprites.com · Personal copy, no sharing or resale · ID RS-20261103-ADA04261

Presidio

- Open source, MIT licence, formerly a Microsoft project, now the community project "Data Privacy Stack". Latest release on 27 October 2026.
- Finds names, emails, phone numbers, IBANs and more, using language models, patterns and checksums. Replace, redact, hash or encrypt. There is also redaction in images.
- **Languages:** the default is English. For other languages you add a language model (spaCy). There are 13 German recognisers (tax ID, ID card, passport, number plates, trade register and more), switched off by default.
- **Limit:** Presidio itself says there is no guarantee it finds everything. Unusual names or indirect hints ("the owner of the bakery on the market square") slip through.
- **Tip:** add your client names from your CRM as your own list, then Presidio finds them reliably.

Presidio · github.com/data-privacy-stack/presidio
`github.com ↗`

`// 04`

Ready-made building blocks

- **LiteLLM Presidio guardrail** (open source): a gateway in front of the model. `pre_call` masks every request, `output_parse_pii: True` puts the original values back into the answer. Claude Code can point at such a gateway via `ANTHROPIC_BASE_URL`.
- **Claude Code hook:** a guard that runs on your own computer before anything goes out. Details in the section "The hook in Claude Code" below.
- **Not recommended:** LLM Guard is archived and has no German. NeMo Guardrails only documents English.

LiteLLM PII guardrail · docs.litellm.ai
`docs.litellm.ai ↗`

Claude Code hooks · code.claude.com
`code.claude.com ↗`

© 2026 REPOSPRITES · @reposprites · reposprites.com · Personal copy, no sharing or resale · ID RS-20261103-ADA04261

// 05

The hook in Claude Code: doesn’t real data already go out?

No, if it's set up right. The key is the difference between two things with almost the same name:

- **Claude Code** is a program that runs on your own computer (or your own server).
- **Claude**, the AI model, runs at the provider: at Anthropic, or via AWS Bedrock in the EU.

A hook is a small script that Claude Code starts **on your computer**, at fixed points.

How it works

1. You type a prompt and press Enter.
2. Before anything goes to the model, Claude Code runs your `UserPromptSubmit` hook, locally.
3. The hook checks the text, for example with Presidio or a list of your client names.
4. If it finds client data, it blocks (exit code 2). According to the Claude Code docs, a blocked prompt **never** reaches the model. You get a message like "Please pseudonymise first".
5. If it finds nothing, the prompt goes to the model as usual.

For files there's a second hook: when Claude reads a file, the reading happens locally. A `PostToolUse` hook can mask the result (`updatedToolOutput`), and only the masked text goes to the model.

Limits, honestly

- **It only stops what it recognises.** A name that Presidio or your list misses goes out. That's why you check samples.
- **Stop, not rewrite:** a `UserPromptSubmit` hook can't pseudonymise your prompt automatically. For that you use the loop or the LiteLLM guardrail.
- **Local only:** in Claude Code cloud sessions (in the browser) your local hooks don't run. The protection applies to Claude Code on your own computer or server.
- **You build it yourself:** I found no ready-made, tested Presidio hook for Claude Code. Test it with sample texts without real data.
- **No replacement for the contract:** everything that isn't stopped goes to the model. You still need the data processing agreement with your provider.

Prompt to copy

```
Build me a UserPromptSubmit hook for Claude Code that runs locally. It should check every prompt with Presidio and my client list clients.txt and block as soon as it finds a name, an IBAN or an email address. The message should say what was found. Add it to .claude/settings.json and test it with three sample prompts without real data.
```

© 2026 REPOSPRITES · @reposprites · reposprites.com · Personal copy, no sharing or resale · ID RS-20261103-ADA04261

// 06

Two honest notes

- **Presidio doesn't catch everything.** Check samples by hand regularly, especially at the start.
- **For you, the data stays personal data.** You hold the mapping list, so the GDPR still applies (Recital 26, EDPB guidelines 01/2025). The loop lowers the risk and helps with data minimisation, but it doesn't replace the agreement with your model provider. Only true anonymisation takes data out of the GDPR.

// 07

Other no-tech providers

PROVIDER	PRICE (ACCORDING TO THE PRICING PAGE)	SPECIAL FEATURE
nele.ai	Credits from €10 net per month	Servers in the EU, Claude available, optional filtering of personal data before the AI
meinGPT	from €18 per user per month plus model budget	Hosted in Germany and the EU, ISO 27001
Mistral Le Chat Team	24.99 USD per user per month	French provider, no Claude

For all of them: ask in writing about backend, region, retention and the data processing agreement.

// 08

Your first hour

- [ ] Decide: simple (Langdock), control (the loop), or both.
- [ ] With Langdock: create a trial and ask the seven questions above in writing.
- [ ] With the loop: install Presidio and switch on the language model you need.
- [ ] Run ten of your own texts without real client data through it and check what it finds.
- [ ] Only then work with real data, with samples.

© 2026 REPOSPRITES · @reposprites · reposprites.com · Personal copy, no sharing or resale · ID RS-20261103-ADA04261

// 09

My tip

Start small with both: Langdock for your team's everyday work, without client data. And a loop with Presidio for the one task that contains names, for example quotes or client emails. That way you get to know both ways before you commit.

All prices, features and links checked on 1 November 2026 on the providers' pages and in the documentation; the security badges on 3 November 2026 on langdock.com/security. Prices change, check before you buy.

— *@reposprites*

Want to make animated videos like this too? Visit my website and get the reel engine.

reposprites.com ↗

© 2026 REPOSPRITES · @reposprites · reposprites.com · Personal copy, no sharing or resale · ID RS-20261103-ADA04261