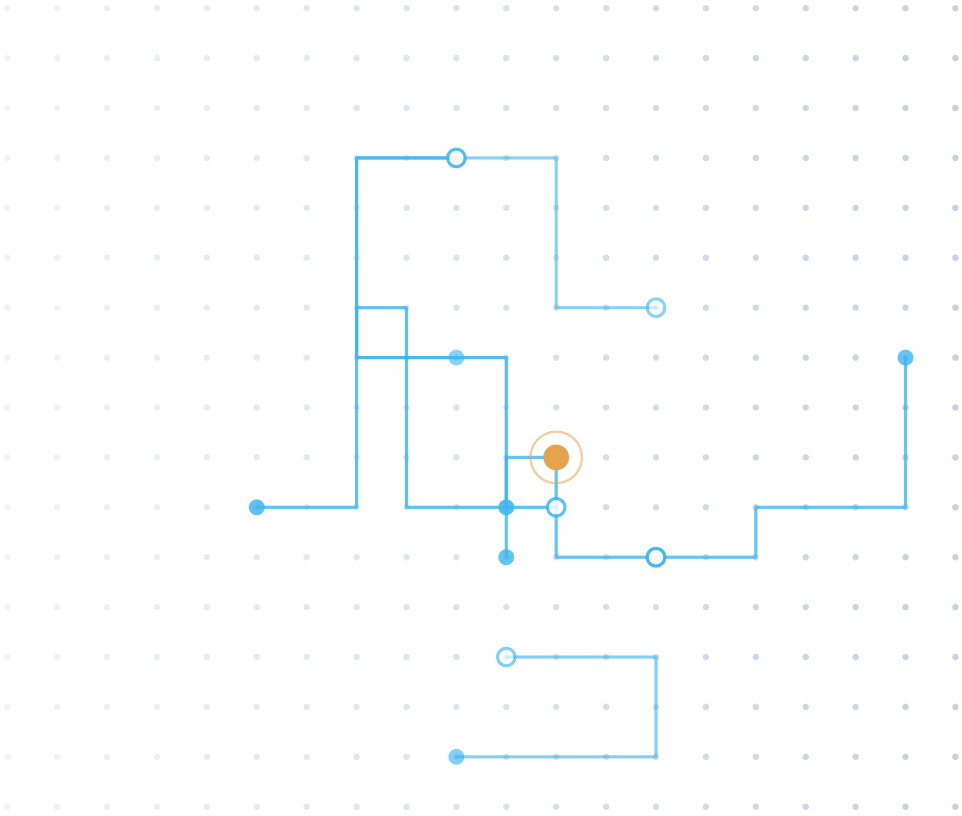


PRACTICAL GUIDE

Privacy checklist for automations that touch customer data

15 checks so your automation uses only the data it needs and nobody sees what they should not.

AUTHOR
Mario Chasco · Flowara AI

FOR
Business owners who handle customer data and the people who build their automations

READING
9 min

EDITION
v2.0 · November 2026

BEFORE YOU START

What you will get from this guide

01 Take a 10-minute inventory of the data a task really uses.

02 Decide what stays on your own machine and what leaves for the cloud anonymised.

03 Separate clients and permissions through structure, not instructions to the model.

04 Arrive at the conversation with your legal adviser with the homework done.

CONTENTS

01 The problem: you hand over more keys than needed

02 A 10-minute data inventory

03 Minimise and anonymise before anything leaves

04 Separate by client and by role

05 What to ask your AI provider

06 Tell people, and write down what you do

07 Example: a gate that fails closed at a school

08 Privacy checklist and what to document

For: Business owners who handle customer data and the people who build their automations · 9 min read

01

The problem: you hand over more keys than needed

When you leave your car at the garage, you hand over the car key, not your house key. With automations we tend to do the opposite. To let AI answer one email we give it the whole inbox; to summarise one record, the whole customer database. It works, until the day something ends up where it should not.
Privacy problems in automations rarely come from a sophisticated attack. They come from three oversights: feeding in data the task did not need, mixing clients in the same place, and trusting that an instruction to the model ("do not mention this") works as a lock. It does not.

That is why I treat privacy as part of the design, not as a promise. Which data goes in, where it lives and who can reach it are decided before the first line is written, and the structure enforces that even when the model gets something wrong.

This guide is for whoever owns the data and whoever builds the automation. It is not legal advice, and it will not tell you whether you comply with GDPR or any other data protection law that applies to you. For that, talk to a professional. What it will do is help you arrive at that conversation with the homework done.

> ✦ **RULE**
>
> The safest data is the data that never enters the automation.

02

A 10-minute data inventory

Before talking about tools, write down what really goes into the automation. Take one specific task, for example confirming appointments over WhatsApp at a clinic, and answer for each piece of data:

DATA	DOES THE TASK NEED IT?	WHERE IT LIVES	WHO SEES IT TODAY
Customer name	Sometimes	Calendar	Front desk
Phone number	Yes, to reply	Business phone	Front desk and owner
Internal notes	No	Notebook or sheet	Owner
Health or children's data	Hardly ever	Records	Person in charge
Payment data	No	Bank or card terminal	Owner and accountant

Three questions complete the inventory. What data arrives without you asking for it? An attachment or an email signature often carries more than it seems. Where does it get copied along the way? Think of logs and backups, and of the sheets in between. How long is it kept?

Everything marked red leaves the automation before you start. Anything amber goes in only when a specific case calls for it.

IN ONE LINE

If you do not know what data goes in, you cannot know what comes out.

03

Minimise and anonymise before anything leaves

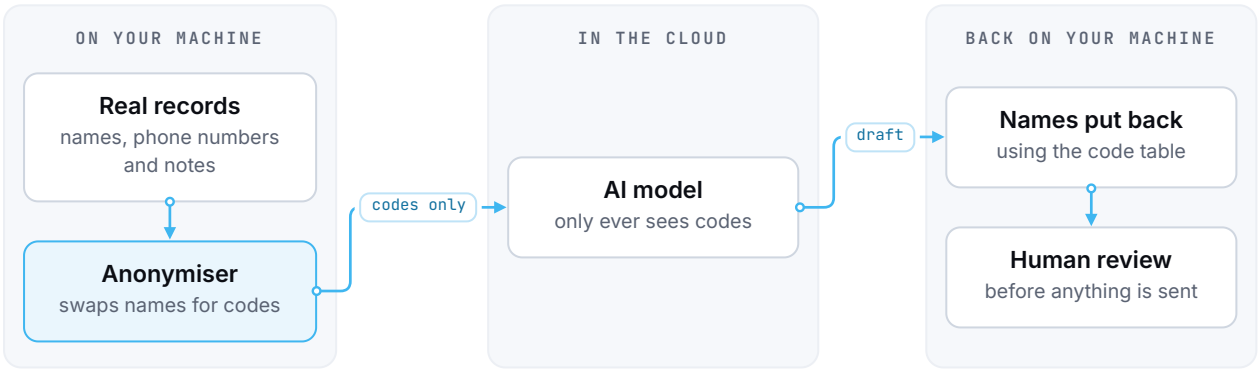
The first two verbs come in this order.

01 **Minimise.** Use only the data the task needs. Confirming an appointment takes a code and a time, not the customer's history.

02 **Anonymise.** If the model has to read a text, every name and identifying detail is swapped for a code before the text reaches it. Do it on a separate copy, never on the real data. And remember that renaming is not anonymising: context can give a person away too.

In practice, this splits the system into two zones. Sensitive data stays on your own machine, and only text with codes goes out to the cloud. When the draft comes back, the names are put back in locally and a person reviews it before it is sent. Health data and children's data do not leave at all, not even anonymised, if the task can be done without them. The table that links each code to a real name lives only on the business's own machine, never in the cloud.

FIGURE 1 **Only codes go to the cloud**



IN ONE LINE

What never leaves your machine cannot leak from somewhere else.

04 Separate by client and by role

The third verb is separate, and here the structure does the work. Every automation is organised in three layers that never mix:

01 **Core:** the shared code, identical for every client.

02 **Client configuration:** opening hours, texts and rules, with no secrets.

03 **Secrets:** keys and passwords, with minimum permissions. They stay out of documents and chats, out of logs and code repositories, and out of any cloud the AI can reach.

And **one container per client**: its own folder, keys and connections. If something fails, it fails inside a single room. The acid test is simple: from inside one client, ask the AI to read another client's files. Any answer other than a refusal is a leak.

> *Privacy is decided by the structure of the system, not by what you ask the model to do.*
>
> FLOWARA AI DESIGN PRINCIPLE

Inside each business, permissions are set per role, before building. An automation acts with the permissions of whoever connected it, not of whoever uses it. If the owner connects it with their own account, anyone who talks to it can reach everything the owner can see.

FIGURE 2 **Who sees what, by role**

	SEES	DOES NOT SEE
FRONT DESK	**Today's calendar** customer name and phone number	**Notes and payments** no internal notes, no takings
OWNER	**The whole business** it is their information	**Keys in plain text** not even the owner
AI AGENT	**Its task's fields** and nothing else	**Other clients** no secrets, no payments
TECHNICAL PROVIDER	**Configuration** and masked logs	**Real data** of end customers, not even in tests

Apply the day-one rule: would you give this access to someone who starts today? If not, do not give it to the AI either. One key per agent, with minimum permissions and read-only until it genuinely needs to write.

IN ONE LINE

Give the AI the access of a new hire, never the owner's.

05

What to ask your AI provider

Every AI service you use processes the data you send it. Before connecting it, ask for written answers to these questions:

01 **Retention:** how long do you keep what I send? Can it be reduced to the minimum?

02 **Region:** in which country is the data processed and stored? Can I choose the European Union?

03 **Training:** do you use my data or my customers' data to train your models? Is that off by default, or do I have to ask?

04 **Sub-processors:** which other companies touch the data, such as hosting, support or monitoring? Do you tell me when the list changes?

05 **Contract:** do you offer a data processing agreement I can sign?

06 **Deletion:** if I leave, how and how quickly do you delete everything?

> ⚠ CAUTION
>
> A well-known provider name answers none of these questions. Terms change with the plan: the free tier and the business tier may treat data differently.

If an answer is unclear, ask again until you have it in one sentence you understand. Keep the answers with a date, and review them once a year or whenever you change plans.

IN ONE LINE

If the provider will not answer in writing, your data does not go there.

06

Tell people, and write down what you do

Anyone who calls or messages your business should know what they are talking to. At the start of a call, the assistant says it is an assistant and whether the call may be recorded. On WhatsApp or in the website chat, the first message says the same.

> **EXAMPLE**
>
> "Hello, you are speaking to the salon's assistant. I can book you in or put you through to a person. Please note this call may be recorded." Have your adviser review the exact wording before it goes live.

On top of that, for every client you build something for, write down on one page what you do with their data:

01 **Which data** the automation handles.

02 **What for**, naming the specific task.

03 **How long** it is kept and when it is deleted.

04 **Which other companies** touch it, starting with the AI provider.

05 **Where it is**: the machine or service, and the country.

That is your record of processing for that client. If the law requires you to keep one, and a professional will tell you whether it does, you are already halfway there. If it does not, the page still helps you answer when a client asks what happens to their data.

IN ONE LINE

What is not written down cannot be reviewed or explained.

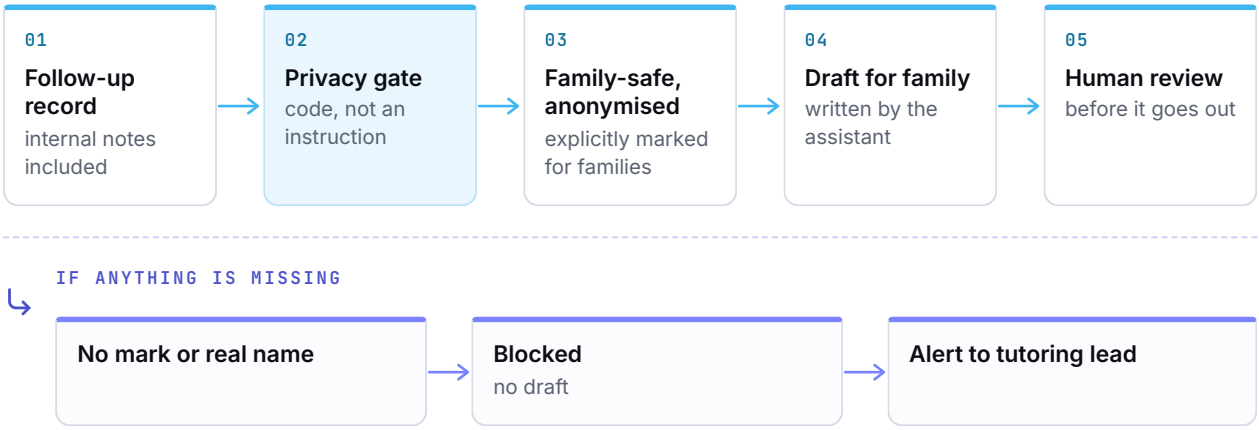
07

Example: a gate that fails closed at a school

A school uses an assistant that drafts reports for families from student follow-up records. Some records contain internal notes that must never leave the building.

The answer is a **privacy gate that fails closed**: a piece of code, not an instruction to the model, that checks every record before the assistant reads it. Only what is explicitly marked as family-safe and anonymised gets through. If the mark is missing, a real name appears where a code should be, or the record cannot be read cleanly, the gate lets nothing through and raises an alert.

FIGURE 3 **A privacy gate that fails closed**



With the synthetic record "Student Alpha, code STU-001", the assistant drafts. If someone pastes a record with a full name, the assistant refuses to draft and explains what is missing.

The order is what matters. The gate sits before the model, so the assistant never gets to read what it should not; nothing depends on the model following an instruction.

IN ONE LINE

When in doubt, the gate closes: better a draft that never goes out than one that leaks.

08 Privacy checklist and what to document

Go through these fifteen checks before you connect an automation to real data:

Before connecting real data

15 ITEMS

- [ ] I have the task's data inventory
  and know where each item lives
- [ ] Every piece of data the task does not need is out
- [ ] Sensitive data stays local or leaves anonymised
  anonymised on a copy
- [ ] Core, configuration and secrets are kept apart
- [ ] Each client has its own container and keys
- [ ] I tried reading one client from another
  and the AI refused
- [ ] No secret is within the AI's reach
  not in documents, chats or repositories
- [ ] Permissions are set per role
  and each agent has a minimal key
- [ ] A privacy gate in code fails closed
- [ ] Logs mask personal data
- [ ] The provider answered in writing
  retention, region, training, sub-processors
- [ ] Calls and messages say an assistant is answering
  and whether it records
- [ ] There is a record of processing per client
  reviewed with a professional
- [ ] Demos and tests use synthetic data only
- [ ] A named person is responsible for privacy

What to document for the client: the data inventory, who can see what by role, the provider's dated answers, the record of processing, how the privacy gate works and who to contact if something goes wrong. A couple of pages is enough, as long as someone non-technical can read them.

NEXT STEP

Free diagnostic (20 min)

If you want to see which part of your business can be automated with the risk under control, book a free 20-minute diagnostic.

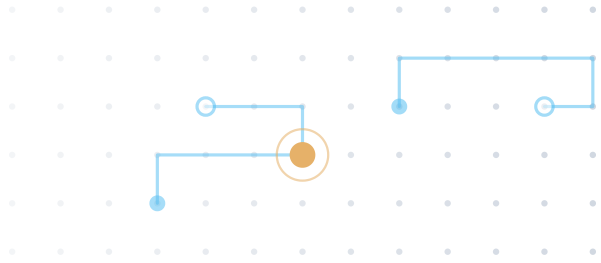
Book a diagnostic ›

flowaraai.com/en/contact

RELATED RESOURCES

Public secretary template with a privacy gate ↗

Real automation cases ↗



Mario Chasco · Flowara AI

AI systems in production for small businesses.

flowaraai.com
github.com/mariochasco78557
youtube.com/@mariochasco_AI
linkedin.com/in/mario-adrian-chasco-frauca

Share this guide as is, crediting flowaraai.com.