The Mirror

Where do you actually stand? Your AI, your data, and EU law.

PiaXel Nexus, EuTrustedIA.eu project

2026-07-16

Contents

Before anything else: this is not an audit **3**

“I don’t do AI” **4**

Step 0 — Does this actually apply to you? **6**
You are established in the European Union 6
You are established outside the European Union 7

Step 1 — The probe: fourteen axes **8**
Axis 1 — What exactly does your system do? 8
Axis 2 — Who are you talking to? 9
Axis 3 — What data do you process? 9
Axis 4 — What model are you using? 10
Axis 5 — Where is your infrastructure hosted? 11
Axis 6 — What legal basis do you have for processing this data? 12
Axis 7 — Does your system decide instead of someone? 13
Axis 8 — How far does your training go? 14
Axis 9 — How many people do you reach? 14
Axis 10 — Where are your users? 15
Axis 11 — Is your use case “high-risk”? 15
Axis 12 — Do the people you interact with know there’s an AI involved? 16
Axis 13 — Where does your contacts’ data come from? 18
Axis 14 — What’s running without you? 18

Step 2 — Where this places you **20**
There is no score, and that’s deliberate 20
And if I raised no flag at all? 21

Step 2bis — What a machine can do for you, and what it never will **22**
What a tool does better than you — because it’s inventory work 22
What no tool will do in your place — and be wary of whoever promises otherwise 23

Step 3 — The framework behind these questions **25**

Now what? Three doors, no obligation **26**

1. Do nothing right now 26
2. See what you've actually deployed 26
3. Move from a finding to proof 26

What this document does not do **28**

Before anything else: this is not an audit

This document is a **mirror**, not a verdict.

It will not tell you whether you are “compliant” — no one can tell you that on your behalf, least of all a PDF. It carries no certification, and we issue none: that is not our trade, and the law does not let us.

What it does is more useful, and more honest: it puts **your actual situation** next to **what EU law says about it**. You answer fourteen questions about what you already do. Against each answer, you read what applies and what it concretely changes for you.

The result is a **declared approximation**. It is only as good as your answers. It replaces neither a data protection officer nor a lawyer — it lets you go and see them already knowing what you’re talking about, which is already considerable.

This document does not constitute legal advice and does not engage the liability of its publisher. It describes general rules; your situation is particular — and that is precisely what separates a text from advice.

Twenty minutes. Something to write with. No one to convince.

> **State of the law as of 16 July 2026.** This document carries its date because the law moves. If you’re reading it more than six months later, the current version — always free — is at eutrustedia.eu/whitepaper. We do not claim to be “always up to date”: we would rather give you a date you can check.

“I don’t do AI”

That is the sentence we hear most often. It is almost always wrong.

You may not have an AI product. You may not train any model, may not write a single line of code. But over the last two years, AI has slipped into the most ordinary gestures of your work. Tick what you **already** do:

❒ **An AI reads, sorts, summarises or drafts your emails** ❒ **You record your meetings** to get an automatic summary ❒ **You hand documents to an AI to analyse** — contracts, CVs, invoices, client files ❒ **You write or post on social media** with the help of an AI ❒ **Your CRM enriches, scores or follows up** your contacts automatically ❒ **You prospect** with lists built, enriched or personalised by an AI ❒ **An AI agent is wired into your messaging apps** (WhatsApp, Telegram, Slack), or drives your computer and your software to run tasks for you ❒ **Automated scenarios** (n8n, Make, Zapier) connect your tools, your databases and your APIs ❒ **You built your website** with a tool that embeds AI

Tick a single box, and this document is for you. Here is why.

What you do	What it means, legally
AI reads your **emails**	You are passing your **clients’** personal data to a vendor, without their knowledge. It isn’t your data: it’s theirs, and you answer for it.
A **meeting summary**	You are processing the data of **every participant** — including the ones who asked for nothing. And what gets said in a meeting (health, conflict, union matters) is rarely trivial.
A **document** handed to an AI	The contract, CV or file you paste into the box belongs to someone. You have just handed it to a third party, often outside the EU, sometimes without knowing whether it will be used to train a model (axis 4).

What you do	What it means, legally
A **post** drafted by AI	Contrary to what you may have been told, the AI Act does **not** require you to label it if you read it back and stand behind it (axis 12). We'll explain why — and who this actually targets.
A **CRM** that enriches on its own	You are collecting data **from somewhere other than the person themselves**. GDPR then requires you to inform them (Article 14). It's probably the most widely ignored obligation in the regulation.
AI-assisted **prospecting**	Same issue, more visible: the person discovers you know things they never told you.
An **agent** or **automated scenario**	Something runs without you, on data you no longer see pass by. It's the heaviest-consequence item on this list — and the least examined (axis 14).

None of these sentences means you're at fault. They mean **you're already in scope** — and the question isn't "does this apply to me?" but "where do I actually stand?"

> **The real trap isn't the technology, it's the ordinariness of it.** No one wakes up thinking "I'm going to transfer health data to a foreign processor without a legal basis." People wake up thinking "I need to write up my meeting notes." Some days, those are exactly the same sentence.

Step 0 — Does this actually apply to you?

Start here. It's the question everyone skips, and it decides everything else.

You are established in the European Union

Yes, GDPR applies. Article 3.1: as soon as processing takes place in the context of the activities of an establishment in the Union, the regulation applies — **regardless of where your servers run**. Your database sits with a US host? That changes nothing about whether you're covered (it changes something else: see axis 5).

> **The myth to kill right now: "under 250 employees, so I'm exempt."** This is false, and it is probably the most expensive belief in this entire document. It comes from a partial reading of **Article 30.5**, which concerns *only* the record of processing activities — and which exempts you from nothing once the processing is **not occasional** (true of any product that runs continuously), carries **a risk** to individuals, or involves **sensitive data**. In other words: for a founder running a service continuously, the exemption almost never applies. And it has **never** exempted you from the rest of the regulation.

If you place on the market or deploy an AI system in the Union, the **AI Act** also applies to you (Article 2). But **not in the same way depending on your role**, and that's the single most useful distinction in the whole regulation:

- **Provider**: you develop an AI system, or place it on the market under your name. You carry the bulk of the obligations.
- **Deployer**: you use, in your own activity, a system built by someone else. Your obligations are **markedly lighter**, and targeted at specific situations.

The vast majority of founders are **deployers** — and believe themselves subject to provider obligations. We'll come back to this at axis 12, where the confusion costs the most.

You are established outside the European Union

Yes, GDPR applies too — and this is what most founders outside the Union don't know.

Article 3.2 doesn't look at where *you* are. It looks at where **the people whose data you process** are. It applies whenever, without being established in the Union, you:

- **offer goods or services** to people located there — paid or **free**;
- or **monitor their behaviour** within the Union (audience measurement, profiling, targeting).

Concretely: your service runs out of Austin, Bangalore, Douala or São Paulo, your company has no European entity, and you have users in Europe. **You are within scope of GDPR.** Billing in a different currency, having only an English-language interface, or never having set foot in Europe changes none of that.

The **AI Act** follows the same logic (Article 2): it covers those who place an AI system on the Union market, **and** those, established in a third country, whose system's output is used within the Union.

> **The obligation almost no one knows about: the representative in the Union.** If GDPR applies to you through Article 3.2, its **Article 27** requires you to **designate, in writing, a representative established in a Member State** — the one where the people whose data you process are located. That representative is the point of contact for supervisory authorities and for data subjects. It isn't a courtesy formality: it's a standalone obligation, and its absence is immediately visible.
>
> **But it doesn't apply to everyone, and we won't saddle you with it without saying so.** GDPR Article 27.2 removes this obligation where your processing is **occasional**, does **not** involve large-scale processing of sensitive data or offence-related data, **and** is unlikely to result in a risk to people's rights. These conditions are **cumulative**: it takes only one to fail for the obligation to come back. An online service that runs continuously with European users is nothing occasional — but a project that served three European customers last year, is.
>
> *(Not to be confused with the **AI Act's own Article 27**, which is something else entirely: the fundamental rights impact assessment. Two regulations, two Article 27s — and we regularly see the two mixed up, including by professionals.)*

> **This isn't a threat, it's a door.** The European single market represents hundreds of millions of people. EU law isn't a wall built against you: it's the price of entry, and it's **the same for everyone** — including your European competitors. Knowing where you stand is knowing whether you get through the door.

Step 1 — The probe: fourteen axes

For each axis: tick what matches your **actual** situation (not the one you're aiming for), then read what follows. When a **flag** goes up, note it — you'll count them at the end.

Axis 1 — What exactly does your system do?

❒ It converses (assistant, conversational agent) ❒ It acts autonomously (chains actions, calls tools) ❒ It sorts or classifies (documents, requests, messages) ❒ It generates content (text, image, audio, video, code) ❒ It recommends ❒ **It scores, rates or ranks people**

What this means for you

The first five boxes describe most projects, and are nothing alarming on their own: it's your answers to the axes that follow that will decide.

The last one deserves a pause. Scoring people is **not inherently prohibited** — banks, insurers and recruiters have done it forever. But two regimes can switch on, and you need to know which one targets you:

- If the scoring leads to people being treated in a way that is **detrimental in a context unrelated to the one in which the data was collected**, or in a **disproportionate** way, you enter the territory of the AI Act's **prohibited practices** (Article 5). This is rare, but absolute: no compliance measure fixes a prohibited practice after the fact.
- If your scoring decides access to employment, credit, education or an essential service, you are probably within **Annex III** — that's axis 11.

> **FLAG RAISED** if you ticked the last box. Not because it's prohibited: because this is the axis that governs all the others most heavily.

Axis 2 — Who are you talking to?

☐ Colleagues, internally ☐ Adult customers, knowingly ☐ **Minors** ☐ The general public, unfiltered ☐ **People in a vulnerable situation** (patients, jobseekers, elderly people, people with disabilities, people in financial difficulty)

What this means for you

The audience shifts the legal regime more reliably than the technology you use.

Minors call for heightened care: GDPR grants them specific protection (their online consent is governed by Article 8, with an age that **varies by Member State** — between 13 and 16). A data protection impact assessment becomes very hard to avoid.

Vulnerable people touch a hard edge of the AI Act: its **Article 5.1.b** prohibits systems that exploit vulnerabilities linked to age, disability or a social or economic situation, in order to materially distort a person's behaviour and cause them harm. The line with "persuasive" marketing is closer than it looks the moment you're addressing people in difficulty.

The trap of this axis: people default to "adult customers" without ever checking. If your service is open on the web with no age gate, your real audience isn't the one you're aiming for — it's the one that signs up.

> **FLAG RAISED** if you ticked "minors" or "people in a vulnerable situation."

Axis 3 — What data do you process?

☐ No personal data ☐ Pseudonymised data ☐ Ordinary personal data (name, email, IP address, customer ID) ☐ **"Special category" data**: health, political opinions, religious or philosophical beliefs, trade union membership, biometric or genetic data, sex life or sexual orientation, racial or ethnic origin

What this means for you

If you ticked the last box, you fall under **GDPR Article 9**, and that isn't a compliance detail: it's **the legal regime of your entire project that changes**. Article 9's default position is a **prohibition** on processing this data, with a narrow list of exceptions — the main one, for a commercial project, being **explicit consent**. A data protection impact assessment (Article 35) stops being optional.

The trap is that you tick this box without knowing it. A few cases where it's the actual topic:

- **your inbox**, once an AI reads it. You don't choose what your customers write to you: a rescheduled appointment "for a chemo session," a delay "because of my depression," an invoice request from a medical practice. You didn't collect this — **you received all of it**, and passed it to your AI vendor along with everything else;
- **your meeting notes**, for the same reason: a mentioned sick leave, a union dispute, a pregnancy announcement. The transcript doesn't filter;
- a conversational agent that books **medical** appointments — you're processing health data, even if all you store is "an appointment reason";
- a nutrition or fitness coaching tool that records weight, conditions or allergies;
- a service that **infers** sensitive data without asking for it: recommending content based on presumed orientation, religion or health status is processing sensitive data. Inferring it is processing it.

> **What makes email and meetings so particular**: in every other case, you decide which data you process. Here, **other people decide for you**. You cannot promise you'll never receive sensitive data — you can only know what happens to your inbox once an AI reads it.

> **FLAG RAISED** if you ticked the last box.

Axis 4 — What model are you using?

☐ A model accessed through an API, operated by a vendor ☐ An open-weights model, self-hosted ☐ **A model I trained**, in whole or in part, on my own data ☐ Several models combined, orchestrated by an agent

What this means for you

This axis decides **who answers for what**, and that's where it gets real.

With a vendor-operated model, you aren't exempt: you remain responsible for what you send it. What you need to check is your **processing agreement** (Article 28) and what the vendor does with your data — in particular whether they reuse it to train their own models. Many consumer-grade contracts allow it; many professional users have never read them.

> **The most common gesture, and the most underrated: pasting a document into the box.**
>
> A contract to summarise. A CV to compare. An invoice to extract. A client file to analyse. The gesture takes three seconds and looks harmless — it's the opposite of harmless.

> That document **does not belong to you**. It contains the name, address, sometimes the health or financial situation of someone who has never heard of your AI vendor. By pasting it, you are **transferring** its content to a third party, often outside the EU (axis 5), under a contract you may not have read, and sometimes to a service that **reserves the right to use it to train its models**. None of this is visible on screen: precisely because nothing appears to happen.
>
> **And here's the part almost no one explains to you.** You may have been reassured that, as long as AI stays your internal tool rather than your product, the AI Act imposes nothing on you. **That's true.** But the AI Act isn't the only text: **GDPR doesn't care whether AI is your tool or your product — it looks at whose data it is.** A purely internal tool can therefore sit outside AI Act scope **and** squarely inside GDPR territory. That is exactly the case here.
>
> The test fits in one sentence: *if the client whose file I just pasted asked me "who did you send my file to, and what happens to it now?" — would I know how to answer?*

If you **train** on personal data, you become responsible for the lawfulness of that corpus — where the data came from, and on what basis you're using it for this purpose.

> **A distinction few documents make, and one that will save you a costly mistake**: the **AI Act's Article 10** (data governance and quality) does **not** apply to every trained model. It sits within the chapter on **high-risk** systems. Training a model does not automatically place you there: it's your **use case** that decides (axis 11). Don't load yourself with obligations that aren't yours — that too, is part of compliance.

> **FLAG RAISED** if you trained on personal data, or if you don't know what your vendor does with your data.

Axis 5 — Where is your infrastructure hosted?

☐ Entirely within the European Union ☐ Within the Union, but with dependencies outside it (content delivery, audience measurement, email, language models…) ☐ Outside the Union ☐ **I don't exactly know**

What this means for you

Two answers are worth stopping on.

"I don't exactly know" is the most common answer, and it's already a useful result. It isn't an admission of incompetence: a modern architecture stitches together ten services in three clicks, and

no one reads "sub-processor" pages. But as long as the answer is "I don't know," you can neither keep your record of processing activities (Article 30), nor answer a business customer who asks where their data lives — and in Europe, that question **always** comes up.

"Outside the Union" is not a fault. GDPR doesn't ban transfers: it **regulates** them (Articles 44 to 49). What it requires is that you know a transfer is happening, on what legal basis, and that you've looked at what local law allows the destination country's authorities to do. What it doesn't forgive is not knowing.

> **A nuance few vendors will tell you**: sovereignty isn't binary. A third country can benefit from an adequacy decision (transfers are then eased, not ignored). A European company can run its workloads on foreign infrastructure. A foreign company can operate a European region — while remaining subject to its home country's law. The useful question is never "is this sovereign?" but **"who can legally access what, and under which law?"** — layer by layer.

> **FLAG RAISED** if you ticked "outside the Union" or "I don't know."

Axis 6 — What legal basis do you have for processing this data?

☐ The person's consent (Article 6.1.a) ☐ Performance of a contract with them (Article 6.1.b) ☐ A legal obligation that binds me (Article 6.1.c) ☐ Legitimate interest, after balancing (Article 6.1.f) ☐ **I've never thought about it**

What this means for you

This is the quietest, most structural axis in the document.

GDPR doesn't ask *whether* you're allowed to process: it holds that **all processing must rest on one of the Article 6 bases**, chosen **before** you start, and not swapped later for whichever one is convenient. Without a legal basis, the processing is unlawful — not "poorly documented," but **unlawful**, from the very first record.

Two costly, very common confusions:

- **"I added a checkbox, so I have consent."** Consent must be freely given, specific, informed and unambiguous. A pre-ticked box, bundled consent for ten purposes, or a service you can't use unless you accept things that aren't necessary: none of that is consent.
- **"Legitimate interest covers everything."** It requires a documented **balancing test** between your interest and the rights of the individuals — a written line of reasoning, one you must be able to show. Without that written record, legitimate interest is an assertion, not a basis.

> **FLAG RAISED** if you ticked the last box — or if you can't name your legal basis, per purpose, in one sentence, right now.

Axis 7 — Does your system decide instead of someone?

☐ No, it informs or assists, a human decides ☐ It proposes a decision that a human **genuinely reviews** (they can say no, and sometimes do) ☐ **It decides on its own**, and that decision produces a legal effect or significantly affects the person (access to a service, application screening, granting or refusal, individual pricing, sanction, moderation)

What this means for you

This is the most counter-intuitive axis in the document, and the one people get wrong most often.

GDPR Article 22 gives every individual the right not to be subject to a decision based **solely** on automated processing that produces legal effects or significantly affects them. The word that costs the most is **"solely."**

Many think a human "in the loop" gets them off the hook. But the Court of Justice of the European Union, in its **SCHUFA** ruling (C-634/21, December 2023), tightened this reading: a human who **systematically rubber-stamps** what the machine proposes does not constitute genuine human intervention. The score produced upstream was itself the decision.

In other words: **if your human reviewer always says yes, you are in a solely-automated decision, whatever your org chart says.** The question isn't "is there a human?" but **"does this human have the power, the time and the information to say no — and do they, sometimes?"**

The reverse is also true, and it's the good news of this axis. Article 22 doesn't condemn automated decisions: it targets those that are made **solely** that way. Genuine human intervention — someone who reviews, who has the authority to decide, and who sometimes overrides the machine's proposal — takes you outside the scope of Article 22. You can perfectly well automate 95% of your processing. What matters is the last step.

> **FLAG RAISED** if you ticked the third box — or if, rereading the second, you realise your human reviewer never actually says no.

Axis 8 — How far does your training go?

☐ I use a pre-trained model, as is ☐ I did light fine-tuning (a handful of examples, instruction adaptation) ☐ **I trained or deeply retrained a model on my own data**

What this means for you

This axis extends axis 4, on one precise point: **what remains inside the model**.

A false idea circulates: "once trained, the model no longer contains personal data, so GDPR no longer applies." The European Data Protection Board examined the question (**Opinion 28/2024**) and its answer is nuanced, but clear on one point: a model is **not automatically anonymous** simply because it was trained. It must be **demonstrated**, accounting for the risk of extraction or re-identification. Until that's demonstrated, the data is presumed present.

The consequence is concrete: if someone exercises their right to erasure, the question "and what about the model?" arises — and "that's technically difficult" is not a legal answer.

> **FLAG RAISED** if you ticked the third box without having documented why your model does not reproduce the data it was trained on.

Axis 9 — How many people do you reach?

☐ A few dozen ☐ A few hundred to a few thousand ☐ Tens of thousands ☐ **Millions**

What this means for you

Scale doesn't create an obligation by itself — but it shifts thresholds.

Article 35 requires a data protection impact assessment when processing is likely to result in a **high risk**, in particular in the case of **large-scale** processing of sensitive data, or systematic monitoring.

> **Honesty on this point**: GDPR sets **no numerical threshold**. "Large scale" is assessed on a bundle of factors — number of people, volume, duration, geographic extent. So don't mistake a number for a rule: ten thousand users on sensitive data weigh more than a million email addresses. This axis is a **marker**, not a verdict — it's the intersection with axes 1, 2, 3 and 7 that speaks.

> **FLAG RAISED** if you reach tens of thousands of people or more, AND you raised a flag at axis 2, 3 or 7.

Axis 10 — Where are your users?

☐ In a single country ☐ Across several EU countries ☐ **In the EU and outside it** ☐ I don't know / I've never looked

What this means for you

This axis mirrors step 0: there you looked at **where you are**, here you look at **where your users are**. It's the second one that triggers EU law.

The moment people located in the EU use your service, GDPR is in play — even if you never sought them out. And if you're established outside the Union, this is where the question of the **EU representative** (GDPR, Article 27) from step 0 comes back.

Symmetrically, if your users are **outside** the EU and you're established there, you don't escape GDPR either (Article 3.1) — and you'll find local law adds on top. EU law is not a master key: **being in order in Europe says nothing about your situation elsewhere.** That's a subject in its own right, and a professional in the target country is the only one who can answer it.

> **FLAG RAISED** if you ticked the last box. Not knowing where your users come from means not knowing which law applies.

Axis 11 — Is your use case "high-risk"?

Does your system serve, even partially, to:

☐ Recruit, screen applications, evaluate or promote people at work ☐ Decide access to education or assess learners ☐ Evaluate creditworthiness, or a right to a benefit or essential service ☐ Intervene in health, safety, or critical infrastructure ☐ Identify people through biometrics ☐ **None of the above**

What this means for you

This is the axis that governs the AI Act most heavily — and the only one in the document where answering "none" is unreserved good news.

If you ticked one of the first five boxes, you are probably within **Annex III**: **high-risk** systems. The regime that opens is **substantial** — data governance, technical documentation, traceability, human oversight, robustness and cybersecurity (Articles 8 to 15). This isn't one more form: it's a way of building.

It's also the axis where **the cost of getting it wrong is the most asymmetric**. Believing you're high-risk when you're not saddles you with heavy obligations for nothing. The reverse is far worse. And the line can be thin: a tool that "helps screen CVs" is in scope; a tool that formats CVs isn't.

What is not enough to place you there: the word "evaluate." Annex III targets uses that **decide** people's access to something important — a job, credit, education, an essential service. A dashboard that scores a campaign's performance, a tool that ranks your tasks, an AI that grades the quality of a text: these are not high-risk systems, even though they "score" things. **Look for the effect on a person, not the verb used.**

> **FLAG RAISED** if you ticked one of the first five boxes. If so, this is the flag to address first — it governs the others.

Axis 12 — Do the people you interact with know there's an AI involved?

First, the question that governs everything else:

☐ **I provide an AI system**: I design it, develop it, or place it on the market under my name ☐ **I use AI systems** built by others, in my own activity

Then, as relevant:

☐ A system **I provide** interacts directly with people (assistant, conversational agent) ☐ A system **I provide** generates synthetic content (text, image, audio, video) ☐ **I publish deepfakes** — image, audio or video content that imitates real people or events ☐ **I publish text intended to inform the public on matters of public interest**, generated by AI and **without** a human review process I stand behind ☐ I use **emotion recognition** or **biometric categorisation**

What this means for you

AI Act Article 50 organises transparency. But it does not place the same obligations on everyone, and this is where the crux sits: **it distinguishes those who provide the system from those who use it.**

If you provide an AI system, two obligations apply to you: designing the system so the person knows they're interacting with AI — unless it's obvious from the context (Article 50.1) — and ensuring the

synthetic content it produces is **marked in a machine-readable format**, to the extent technically feasible (Article 50.2).

If you use systems built by others, the article only targets specific cases: **deployers must disclose emotion recognition or biometric categorisation systems** to the people exposed to them (Article 50.3), and **deployers must disclose deepfakes** and **text published to inform the public on matters of public interest** (Article 50.4).

> **And now, what almost no one will tell you, because fear sells better.**
>
> You may have been told that soon you'll need to "label all AI-generated content," your posts included. **Read the article.** If you draft a marketing message with AI's help and read it back before publishing it:
>
> - the **machine-readable labelling** obligation (50.2) sits with the system's **provider**, not with you;
> - the **deployer's** obligation (50.4) only covers deepfakes and text that **informs the public on matters of public interest** — a marketing message is neither;
> - and the text **explicitly** provides that the obligation lifts where the content has undergone **human review or editorial control**, with someone taking responsibility for it. **If you read back what you publish, you're the one taking responsibility for it.**
>
> This document will not load you with an obligation that isn't yours. There's plenty enough to do with the real ones.

Article 50's obligations become applicable on 2 August 2026. That date is tied to **this article, and only this article**: the AI Act does not have "a" single deadline, but a staggered timetable where each part has its own (Article 113). Be wary of anyone who talks to you about "the AI Act deadline" in the singular — it's the first sign you're being sold anxiety rather than law.

> **FLAG RAISED** if you provide a system that interacts with people or generates synthetic content **without either being announced or labelled** — or if you distribute **deepfakes** without disclosing them.
>
> **No flag** if you use AI to draft content that you read back and stand behind. That's the situation of most people reading these pages.

Axis 13 — Where does your contacts' data come from?

☐ They gave it to me themselves (form, sign-up, purchase, direct exchange) ☐ **I found it elsewhere** (professional networks, directories, purchased lists, websites) ☐ **A tool enriches it automatically** (role, headcount, technologies, contact details, score) ☐ I don't know how my database came together

What this means for you

This is the most widely ignored obligation in GDPR — and it hits prospecting and modern CRMs squarely.

When you obtain data **from somewhere other than the person themselves**, **Article 14** requires you to **inform them**: who you are, what you do with their data, on what basis, for how long, **where it came from**, and what their rights are. This isn't triggered by a complaint: it's a **spontaneous** obligation.

And the timeline is precise: within a reasonable period, **at the latest one month** after obtaining it — or, if you use it to contact them, **at the latest at the time of your first message.**

In other words: **your very first prospecting email must itself carry this information.** Almost none do.

The excuse that doesn't hold. Article 14.5.b provides an exception when informing would take "disproportionate effort." It gets invoked wrongly, constantly: sending an email to someone whose address you already have — since you're writing to them — has never amounted to disproportionate effort.

> **The simplest test in this whole document.** Reread your last prospecting message and ask yourself: *if this person asked me "how did you get my details, and what else do you know about me?" — is my answer already in my message?* If not, that's the gap. And it isn't a lawyer's nicety: it's exactly the question people ask, with the very concrete feeling of having been watched without knowing it.

> **FLAG RAISED** if you ticked the second, third, or last box without your contacts having been informed.

Axis 14 — What's running without you?

☐ Nothing: I run every tool myself, I see what it does ☐ **Automated scenarios** (n8n, Make, Zapier…) connect my tools, my databases and my APIs ☐ **An AI agent is wired into my messaging apps** (What-

sApp, Telegram, Slack, email) and reads what gets said there ☐ **An AI agent executes actions** on my behalf: browses, buys, fills in forms, sends, posts ☐ **An AI agent drives my computer** or my software ☐ I don't exactly know what my automations pass along

What this means for you

This is the last axis, and probably the heaviest — because it's the only one where **you're no longer in the loop at the moment it happens.**

An agent wired into your messaging apps reads other people's messages. Your customers write to you on WhatsApp, your partners on Slack, your prospects by email. None of them consented to an AI reading their conversation, and none of them know. You installed a tool — but the data flowing through it isn't yours. Same issue as axis 3, worse: on a messaging app, **you choose neither the content, nor the timing, nor the sender.**

An automated scenario moves your data around without you seeing it leave. An n8n or Make chain is a **series of transfers**: your database to an AI, the AI to your CRM, the CRM to a messaging service. Every arrow in the diagram is a transfer under GDPR, and every service has its own jurisdiction (axis 5). The scenario itself never pauses to ask whether that's actually lawful. **A diagram that fits on one screen can contain five transfers outside the EU** — and it's precisely because it's visual and simple that no one counts them.

An agent that executes actions raises one more question: what exactly did it do? GDPR doesn't only ask you to be in order, it asks you to be able to **demonstrate it** (Article 5.2). In other words: know what was done, on what data, when. An autonomous agent with no trail doesn't just leave you exposed the day of an audit — it stops you answering the person concerned who simply asks what you know about them. And if it **decides** something that affects someone, you're back at axis 7.

An agent that drives your machine sees everything your machine sees. Not just what it needs. Yet GDPR requires that the data processed be **limited to what is necessary** (Article 5.1.c). An agent with access to your entire file system to draft a quote processes a thousand times more data than necessary — and you can't say which.

> **The question that sums up this axis**: *if I were asked for the full list of everything my automations read, sent, and to whom, this week — could I produce it?*
>
> Almost no one can. It's not a moral failing: it's the normal consequence of tools built to be forgotten. But **what's been forgotten keeps running**, and you're the one who answers for it.

> **FLAG RAISED** if you ticked anything other than the first box and you cannot say, today, which data passes through your automations and to which countries.

Step 2 — Where this places you

Count your flags. Note their numbers — not just the total.

My flags: _____ in total, on axes number _____

There is no score, and that's deliberate

We will not give you a score out of 100. Not out of caution: **because an average score lies.**

A project that raises a single flag at axis 11 (high-risk) sits in a far heavier situation than a project raising five flags on secondary axes. Merging the two into one number would produce a comfortable, false figure. In compliance, what matters is never the average: **it's the weakest link.**

So read your flags this way:

Three flags outrank the others. If they're raised, they come before everything else:

Axis	Why it governs
Axis 11 — high-risk	It opens an entire regime (AI Act Articles 8 to 15). It shapes how every other axis should be read.
Axis 3 — sensitive data	It changes the legal regime of the whole project, not just one part of it.
Axis 7 — automated decision	It creates enforceable rights for individuals, exercisable immediately.

The other flags are read together, never in isolation. A flag on axis 5 (unknown hosting) is uncomfortable. The same flag, combined with axis 3 (sensitive data) and axis 9 (large scale), describes an entirely different situation. **It's the intersections that speak** — and that is precisely what a document on paper cannot do in your place.

And a flag on axis 14 multiplies all the others. What you do by hand, you do once and you saw it happen. What an automation does, it does **a thousand times, with no witness.** A sensitive document

sent by mistake to a vendor outside the EU is an incident; the same gesture, wired into a scenario that runs every night, is a **pattern.** If you raised a flag on axis 14 **and** another one elsewhere, that's the intersection to look at first.

And if I raised no flag at all?

Two possibilities, and both deserve honesty.

Either your project genuinely is simple and well-scoped — that exists, and it's excellent news.

Or you answered what you **think** you do, rather than what you **actually** do. That's the most common answer, and there's nothing shameful about it: no one knows by heart the list of services running inside their own tools. That's exactly the subject of axis 5, and why "I don't know" is a legitimate answer there.

> **The only result with no value is the one you get by reassuring yourself.** If you hesitate on an axis, raise the flag: a noted doubt is worth more than a box ticked for comfort.

Step 2bis — What a machine can do for you, and what it never will

You have your flags. The next question is fair: **what can be tooled, and what stays your own work?**

We answer honestly, including when the answer doesn't suit us. A tool vendor who claims to automate everything is lying to you; the worst service we could do you is let you believe a piece of software will decide in your place.

What a tool does better than you — because it's inventory work

Your flag	What a machine can do
Axis 5 — "I don't know where my data is"	**Draw the map.** A tool can trace back the services, models and databases you call, and attach each one to its jurisdiction. That's inventory work: long, thankless, judgment-free — exactly what a machine does better than a tired human.
Axis 4 — "I don't know what my vendor does with my data"	**Read the terms for you.** An automated audit goes and finds what a service says publicly about its data handling, its transfers and its training practices. You will never read those pages. A tool will.
Axis 4 — "I paste documents full of personal data"	**Strip personal data before it's sent.** Automated document anonymisation is a technical problem, and tools solve it. The risky gesture (pasting the raw contract) has a tooled alternative.

Your flag	What a machine can do
Axis 14 — “I don’t know what my automations pass along”	**Read the scenario and count the arrows.** An n8n or Make export is a file: a machine can parse it and tell you which services it crosses, and to which countries. You just see a tidy diagram.
Axis 14 — “my agent runs with no witness”	**Keep the trail.** Observing what an agent did, when, on what — that’s precisely what a human cannot do continuously, and what a log does with no effort at all.
Axes 2 · 3 · 9 · 11 — impact assessment, high-risk systems	**Pre-fill the documents.** An impact assessment, a record of processing, a transparency notice: their structure is known, and 80% of their content can be derived from your map. Copying it out isn’t the same as thinking it through.

What no tool will do in your place — and be wary of whoever promises otherwise

Your flag	Why it’s irreducibly human
Axis 6 — the legal basis	**No one can choose your legal basis for you.** It’s a legal judgement about *your* purpose, your relationship with the person, your balance of interests. Software that “assigns” you a legal basis is making it up. We won’t do that.
Axis 7 — the human who says no	A tool can detect that your decision is automated. It cannot **be** the human who contests it. That authority doesn’t install.
Axis 11 — the high-risk classification	A tool can say “your case resembles Annex III.” **The classification itself has to be owned** — and discussed with a professional, because it carries consequences.

Your flag	Why it's irreducibly human
Axis 13 — informing your contacts	We can remind you that you owe it to them. Writing to your prospects is you.
All axes — the signature	**A document produced by a tool is only worth what the person who reads and signs it is worth.** A data protection officer, a lawyer. We document; **we certify nothing**, and no honest party will either.

> **What this means, concretely.** Read the two tables together: the machine takes the **inventory** (what exists, where, with whom, what has moved); you and your advisers keep the **judgement** (why, on what basis, with what responsibility). Any pitch that offers you the reverse — software that judges, or an expert who does three billed days of manual inventory — makes you pay for the wrong side of the line.

> **Something you may have just noticed.** The questions in this document aren't those of a standard compliance checklist. They almost never ask *"do you have document X?"* — they ask *"do you know what's happening inside your own operation?"* That's not a coincidence: **it's the same conviction that produced this document and the tools we build.** You can't be in order about what you can't see. Everything starts with seeing.

Step 3 — The framework behind these questions

These fourteen axes didn't come from nowhere. They follow from five principles we apply to ourselves before proposing them to anyone else. Here they are, one sentence each — the full reasoning lives at eutrustedia.eu/eudai.

1. Radical sovereignty. Know where your data is, with whom, under which law — layer by layer, never through a single global label. The useful question isn't "is this sovereign?" but "who can legally access what?"

2. Self-demonstrable compliance. Compliance isn't a state you declare, it's work you show. What can't be demonstrated doesn't exist on the day someone asks.

3. Human in the loop — genuinely. A human who always approves is not a human in the loop (axis 7). The question isn't their presence, it's their power to say no.

4. Transparency of training data. What goes into a model decides what comes out of it — and what you can answer for.

5. Output transparency, input robustness. An AI is also judged by what it lets in, not only by what it produces.

Now what? Three doors, no obligation

You have your flags. Here is honestly what you can do with them — including nothing.

1. Do nothing right now

That's an option, and sometimes the right one.

If you're still validating your idea, if your product has no users yet, if you raised none of the three governing flags: **you've just gained the essential thing.** You now know where to look the day it moves, and you won't spend ten months building in the wrong direction. Keep these pages, redo the probe once your project has changed.

No one will follow up with you. This document asked you for nothing.

2. See what you've actually deployed

If axis 5 left you uneasy — "I don't exactly know" — that's the most useful signal in this whole document, and it has a free next step.

At eutrustedia.eu/wizard, declare the building blocks of your system (models, hosting, databases, third-party services): you get a map of your stack, each block classified by jurisdiction, and your points of concern. No account needed, a few minutes, and your blocks stay in your browser.

It's the natural next step from the mirror: here you looked at **your situation**; there you look at **your system**.

3. Move from a finding to proof

This document tells you **where you stand**. It doesn't tell you what to do, in what order, and it leaves you nothing to show — and that's deliberate: these are two different jobs.

The day you need a prioritised roadmap, enforceable documents (impact assessment, processing record, transparency notice) and dated proof of your diligence, that's where our offer begins: eutrustedia.eu/pricing.

And if your situation calls for a human eye — a data protection officer, a lawyer, an expert — our directory lists them: eutrustedia.eu/experts. We sign nothing on their behalf, and we take no commission on their work.

What this document does not do

In the interest of honesty, and because you deserve to know it before drawing conclusions:

- **It does not declare you compliant.** No one can, least of all us. This isn't editorial caution: we document compliance, we don't certify it — that's a distinct, regulated profession, and not ours.
- **It replaces neither a data protection officer nor a lawyer.** It lets you go and see them already knowing what you're talking about. That's already a great deal: most people who consult them spend the first hour reconstructing what you just did in twenty minutes.
- **It does not constitute proof.** A document you filled in yourself is not enforceable against anyone.
- **It only covers European Union law.** If you sell elsewhere, other rules stack on top — and being in order in Europe says nothing about your situation in the United States, China, or anywhere else. We do not claim to know those.
- **It will age — and we'd rather tell you than let you find out.** This document reflects the state of the law on the date printed on its cover. The law moves: case law arrives, the European Data Protection Board publishes guidelines, the AI Act's timetable rolls out in stages. **If you're reading these pages more than six months after that date, go find the current version at** eutrustedia.eu/whitepaper — it's always online, always free, and carries its own date.

  We don't promise you an "always up to date" document: that would be unverifiable, and therefore worthless. We promise you a **date**, one you can check, and a document that **accepts it will age**. For the substance that moves — the texts, the deadlines, the interpretations — the web is more honest than a PDF: the living version lives at eutrustedia.eu/eudai. These pages keep what moves the least: **the questions.**

What it does do, though, no one else will do for you: look you in the eye for twenty minutes.

EuTrustedIA.eu — a PiaXel Nexus project.