GDPR

A simple guide.

Written by Miloš Novović. Feel free to share/copy/ritually burn. Personal opinions, not legal advice. Obviously.

Hey, you!

> *Yes, you!*

Hi there!

You know how sometimes you read one of those crappy things like *"people decide if they wanna continue reading something within 2 seconds"* and you have *zero* idea if it's actually true, but it *kinda* sounds right, so you believe it, and then it comes to haunt you at the most inopportune of moments?

Well, I *do* feel *very* haunted by that particular piece of info please, so *please* give me... 4.3 more seconds.

Wait, did I just use one?

Damn. OK.

Real fast.

What is this?

It's a guide to the GDPR.
Written in a (somewhat) human language (thanks, EU).
Aimed at developers, but hopefully pretty accessible to anyone.

Why?

Because the text of the GDPR is about as easy to read as a hex dump scribbled down by a monkey. A drunk monkey.
Somebody needs to make the rules semi-coherent and, well, practical.
I, humble as I am, decided that I know better than the entirety of the EU legislature, and decided to do it.
Well, at least try.

Who?

I'm Miloš Novović, an Associate Professor of Law, working with law and tech. I'm in Oslo, Norway. (No, I do not enjoy skiing.)
I'll drop the mandatory 72 paragraphs on the depths of my infinite wisdom. You can find the said paragraphs on my homepage (**www.milos.no**).
You'll get to know me a bit along the way, and I'll get to know you if you drop me a line!!
More on that in the end.

Still with me?

Awesome!

Let's talk *goals*.

How?

So what's the plan?

It's simple. I want to cover *all the basics* of Article 5/6, which, in my opinion, is *precisely* **95.671% of the GDPR compliance.**

I am allergic to charts/infographics/fancy PowerPoints which convey no substance. But *here*, let me practice my consulting skills:

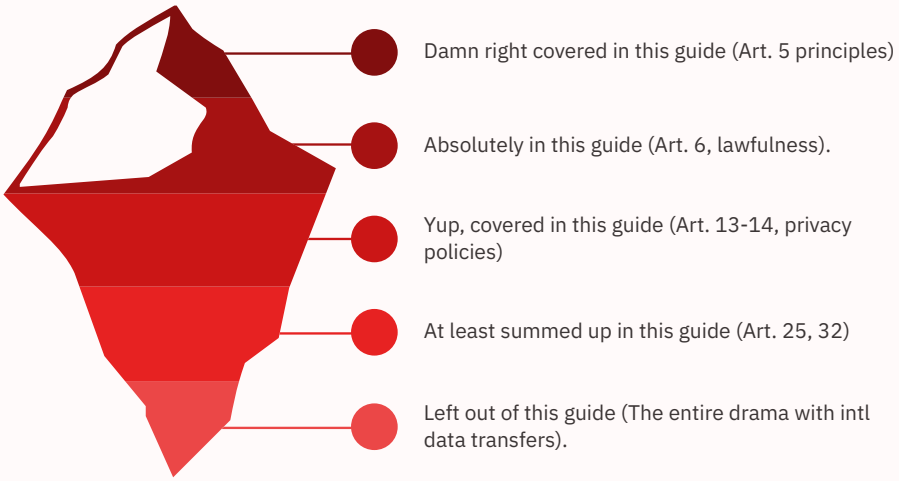

(*Wait,* is this how you do it? Because usually, like, lore videos have the most obscure stuff on the bottom, but then again, there's the expression *'that's just the tip of the iceberg'* which means that the real, systemic stuff is underneath? So is the most important stuff on the top or the bottom??)

See, *that's* why I hate infographics.

The point is: I will teach you the ***important*** stuff.

The building blocks. Equivalent to teaching a new Swift developer about reference vs value types, primitives, functions, basic SwiftUI, MVVM: it will let them build tons of stuff. Sure, a day *shall* come when they will have to research an arcane API, or incorporate a feature which is important, but the lack of which does not *break* the app.[1] But that's a battle for another day.

If you are one of those '*every-letter-of-the-GDPR-is-equally-important-blah-blah-fundamental-right*s'-people, well, you go ahead and write your own guide. I'm sure it'll take the world by a storm. Let me know.

That reminds me. Naturally, obviously, *clearly,* nothing here is legal advice. It's a *generic* legal interpretation made by a guy on the Internet who had too much time on his hands. It's for free. I don't advise for free. I only preach for free. Deal? Deal. Awesome!

Summary

So, what's *really* in here?

Everything, here, all at once

"I don't have the time to read all of this! TELL ME WHAT I NEED TO DO."

OK. Here is *everything* we are going to cover in this guide. Foolishly enough, as a combo of a ToC *and* a TL:DR.

It might look intimidating right now, but! We're gonna do it one step at a time, and I promise, it won't be horrible. So.

> **Chapter 1** establishes *whether the GDPR applies to you or not.* First, we talk territorial scope — what if you are an American company, an indie dev from Argentina? Then, we talk material scope — when does it kick in? What's 'personal data'? If your app just uses Firebase for authentication, and you have zero way of actually finding out who a user is — does it still apply? (Spoiler alert: yes.) How can you anonymize? (You likely can't.) There is some cosmic injustice in the first chapter having to be the most boring one, but ummm, skip if you're sure you know this stuff.
>
> **Chapter 2** gets you up and running — *what do you need to do first?* It's actually simple. You will need to sit down and make a list of *purposes* for which you are using people's data. ("To let people sign up for my newsletter, to optimize the connection speed, and to debug my app.")
> The chapter then follows up with the main rules of the GDPR — the central one of which is '*if something serves no purpose, delete it*'.
> Seriously. It's *that* simple. But it's a *huge* chunk of what the GDPR demands.
> We will also talk about how secure the data has to be (depends on the context; a horoscope app != a health-tracking one).
> *Don't* skip this one. It's the most important part of the guide.
>
> **Chapter 3** is for those who wonder about *consent.* Do you need to ask users for it? When/how?
> The spoiler, I guess, could be that *no*, you don't always need to ask for consent, and *yes*, the GDPR is actually reasonable enough to say 'if your app can't work without this data, don't bother with consent'. Actually, it pretty much says the same about anything that is not creepy, that serves a clear goal, and that a normal person would not find objectionable. In other words — it lets you make a lot of stuff opt-out. Which is both nice and reasonable.
> We'll learn about the alternatives, and then, for each purpose we jotted down in Chapter 2, figure out what to do.
>
> **Chapter 4** is about *privacy policies.* Why? Because everyone needs one, and because everyone talks about them. *No*, they do not have to be long, or boring, or detailed. You take your list from Chapter 2 and 3, and you add a tiny bit of extra context. No need for lawyers. No need for AI. Just common sense. (And a bit of refactoring.)
>
> **Chapter 5** is ummm a... kitchen sink. First, will I try to cram in some info about the things I didn't cover in this guide, even as a single line. (What happens if you experience a data breach, say.) Secondly, we talk about the fines and how to avoid getting swindled by those selling 'GDPR compliance'. Lastly, we wrap up with some platform-specific considerations (because the GDPR is just one part of the picture, and companies like Apple have their own requirements).
> It's not meant to be comprehensive. It's basically just one giant P.S.

My recommendation? Take it one chapter per day.

Lies, damn lies!

Before we start: there's a lot of *wrong* info on the GDPR out there. Which I have Strong Opinions™ about. Which most people won't need/want to hear, but I think they might be important to get out of the way before we dive in.

Feel free to skip. But some might surprise you!

Mythbusting!

GDPR == fines

Nah, not *really*.
You know the scary number that's being tossed around — 4% of global turnover, 20 million EUR? It is very, very, *very* unlikely you will be slapped with something even *remotely high*, unless you are a major company doing some shady stuff. Important limitations apply — legally and practically.
First, fines have to be *proportionate.* Secondly, there are limits to the capacities of the enforcement authorities, and they typically do not prioritize cases where someone forgets to update their privacy policy subsection 8 when switching vendors. They *might*. In the sense that you *might* be ran over by a tractor driven by a chicken and a giraffe gangster duo tonight.
Which does not give you a free pass, because, as we will see, all the stuff in here is pretty much reasonable!

GDPR == a perpetual 'no':

"Yeah, I'd like to do that, but I can't. You know. The GDPR."
Heard that one before?
I say: if what you are doing is *sensible*, the GDPR will *absolutely* let you do it. It might ask you to document the living hell out of it and to plan ahead, but that's *fine*.
And yes! If you can prove me wrong, send me an e-mail, and you have a free beer in Oslo.
Or a mocktail. Your call.

GDPR == consent:

Nope. As mentioned earlier and explored *ad nauseam* later, nope.[1]

GDPR == access/deletion:

Yup, people have rights. Nope, you *do not* need to delete the data whenever they ask for it or give them full access to it. Reasonable checks apply. Don't jump the gun.

GDPR == crazy security:

No, see, IT security and security under the GDPR are not the same. The GDPR does not ask you to enforce a *uniform* security level — it simply states that you need to look at the type of data you have, risks it poses, and try to protect it to a reasonable extent. A data breach is, therefore, not automatically a GDPR violation. A grocery store hacked by a nation state has not violated the GDPR — they had no duty to mitigate against the risks which are so unforeseeable. The rule is simply — do what a good dev in your field would do, bearing the specific use-case in mind. No plaintext passwords. 2FA where possible/reasonable. The works.

GDPR == hard:

No. It just takes some time to wrap your head around. As I said — most of the answers are pretty much 'common-sense-meets-good-coding-practices'; the issue is that it is tough to know which *questions* to ask.
But that's why this guide is here!

Let's go!

01

Does it apply?

Does the GDPR apply to me?

> *Most likely, yes.*

Each law has what lawyers like to call 'scope of application', which can be 'material' or 'territorial'. That is just fancy speak for **'we need to know what the law regulates and in where in the world it applies'**. If I inherit something in Norway, Norwegian inheritance law applies. If I rob a bank in Greece, Greek criminal law will apply. And so forth. Typically.

But things can get complicated. What if I gave Snow White a poisonous apple in Germany, which she ate in France, only to fall asleep in Spain, which lead to me inheriting her castle in Germany? [1]

It sounds like a bad joke, but it *is* relevant: if you are an indie dev from Brazil, and you process nothing more than hashed usernames of people in the EU, do you have to worry about the GDPR? What if you are a medium-sized SaaS company registered in the EU and you make a todo app which sends you zero data, but has iCloud sync enabled? Do you need a privacy policy, consent forms, risk analysis, anything GDPR-related?

Well, let's see!

01.1

Where?

Where does the GDPR apply?

> *In the entire Universe, technically.*

The starting point for most laws is simple: they apply in the countries where they were adopted (makes democratic sense). But some countries decide that some of their laws should apply abroad as well; that's called 'extraterritorial application'.

The GDPR is one of those.

First, it applies to anyone with an '**establishment**' in the EU. That doesn't necessarily mean registered offices, just some stable presence — your freelancers are in the EU and are doing A/B testing with EU users, for example. Some edge cases. But let's keep it simple for the time being: **EU-company == GDPR.**

Now, *without* an active EU establishment, the GDPR applies *only* if you're '**offering goods or services**' to people in the EU or '**monitoring**' their behaviour.[1]

'Goods and services': it's enough to show you're trying to reach customers/users in the EU – whether your app is free makes *no difference*. Nor does it matter if you have zero or five million downloads. The question the courts ask is simple: is it obvious, from all the circumstances, that you *intended* to offer something on the EU market? Can people buy or download your app for free? Leave reviews?

For most people reading this, the answer is likely 'yes'.

Alternatively, the GDPR applies if you're 'monitoring' behavior in the EU tracking or analyzing people even without offering anything. Think 'analytics cookie on a website'. Or an alien species trying to establish contact with us. Whatever works.

Examples
Is your free app on the EU App Store? You're offering services == GDPR applies.
Can your site be, well, *technically* accessed from the EU, but people there can't buy anything, get support, leave reviews, and are not being tracked? Then no, the GDPR doesn't apply.
You don't want to offer your app in the EU, but a person who downloaded it Mexico travels to Germany? Nope. You never intended to offer the service in Germany. Forget the GDPR.

The GDPR does not apply to 'data of European citizens'.

It has nothing to do with citizenship, actually.

Or the physical location of your servers.

Your *intention* to offer the service in the EU or monitor the people is the only thing that matters.

Fines outside of the EU

> *Extraterritorial enforcement is lacking.*

How likely is someone *outside* of the EU to be fined?

It's a tough question.

Trying to enforce your laws in other countries is like telling overprotective parents how their kids should behave. It requires a bunch of resources, so for the most part, authorities tend to stick to investigating bigger companies, or those dealing with super sensitive data. To date, there have not been that many cases against non-EU entities, and we have seen barely a handful against SMEs.

But also – and I will come back to this – *please* stop thinking of the GDPR in term of the fines. Fearmongering sucks. People who try to sell it to you suck double.

A small table of notable extraterritorial cases which I am currently aware of is under.

Company	Alleged violation	Status
Grindr LLC	Shared GPS & sexual orientation data	NOK 65 mil; fine upheld on appeal
OpenAI	Trained ChatGPT on personal data	EUR 15 mil; under appeal
Uber Technologies	Sent EU drivers' sensitive data to US servers	E290 mil; under appeal
Locatefamily.com	Had no EU representative; ignored erasure requests	EUR 525 K + EUR 20 K/day; unclear

Don't wanna pay?

> *Well... You probably should.*

What if someone outside the EU gets a fine, and they do not want to pay?

Generally, it's a *horrible* idea.

Customers do not like it. App Stores most certainly do not like it. Business partners avoid it like a plague.

Naturally, most companies therefore chose to comply voluntarily and to file an appeal to the courts of the country where the decision was made. But forcing someone outside of the EU to pay is still *tough*.

Even the authorities say so.

▼ **Official EDPB study**

> "Enforcement of EU SAs' decisions in courts in California and the UK may prove difficult-if not impossiblein a reasonable timeframe without prejudice to its financial cost. However, the adoption of the California Consumer Privacy Act of 2018 may open the door to active cooperation with the California Privacy Protection Agency. Similarly, cooperation with the UK Data Protection Commissioner seems possible in the context of Treaty 108+ (Articles 16 and 17), combined with the functions and missions vested in the UK Commissioner by the UK GDPR and Data Protection Act (DPA 2018). It is worth noticing that the UK Commissioner has concluded a relatively high number of Memorandum of Understanding (MoUs) with foreign data protection authorities. Cooperation with China seems possible in theory but would require a comprehensive approach, including close cooperation with the EU and the Member State public authorities responsible for that relationship with China."
>
> European Data Protection Board, Study on the enforcement of GDPR obligations against entities established outside the EEA but falling under Article 3(2) GDPR, 2021

01.2

When? Why? How?

Scoping, scoping.

Yaay, we are done with the 'territorial scope'!

Which means you can now impress people at parties.

No, seriously, tell them that '*the territorial scope of the GDPR extends to activities of non-EU establishments insofar as there is an intention to target data subjects via offering of goods or services or monitoring their behavior*', which is a *simplified* version. A great conversation starter!

What is an even better one?

'*Material scope*', clearly.

In normal language: **which data** triggers the application of the GDPR? Say you have no idea who a person is — you merely have their randomly generated user id. Or, say, that user id is not even accessible by you — an SDK you are using generates it on a remote server. *Surely* you can't be liable for that?

Well...

I have some bad, then good, then bad, then meh news for you.

Ready??

When does the GDPR apply?

> *Like, always.*

What does the GDPR regulate? When do I have to follow it? How can I avoid it? What if I only have a freaking user session id?

The answer seems simple: the GDPR applies when you are 'involved' in 'processing' of 'personal data', we're told. Ok, next.

Well, no. Not *so* fast. As with all-things-legal, we have to do it one step at a time.

We will do this *a* lot from now on — take it slowly and examine each step separately.

I don't want to teach you legal methodology, because, in my experiences, those classes are all like 'teleological interpretation' this, 'ancient philosophy' that. But the basics of working with law — any law — are important. At the very core, it requires you to take a rule, break it down into its small elements, and then re-assemble it again.

Think of it like code. All things have to fall in place. All vars must be passed in an initializer — or the compiler is going to scream to no end. Or worse yet, it's gonna compile, but you'll *so* wish it didn't. Even where the defaults are... Sensible.

```
enum DataType { case personal, anonymous }

struct GDPR {
  var dataProcessing: Bool
  var dataType: DataType
  var amSomehowInvolved: Bool

  var isApplicable: Bool {
    dataProcessing == true &&
    dataType == .personal &&
    amSomehowInvolved == true
  }
}
```

Take-away: when I mention that applicability calls for x, y and z, and that there are four criteria for consent, you have to examine *each*.

Not ideal, but we're gonna make it!

Sources

GDPR Articles

ArticloA Definitions

For the purposes of this Regulation:

(1) 'personal data' means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person;

(2) 'processing' means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction:

Recitals

(26) The principles of data protection should apply to any information concerning an identified or identifiable natural person. Personal data which have undergone pseudonymisation, which

could be attributed to a natural person by the use of additional information should be considered to be information on an identifiable natural person. To determine whether a natural person is identifiable, account should be taken of all the means reasonably likely to be used, such as singling out, either by the controller or by another person to identify the natural person directly or indirectly. To ascertain whether means are reasonably likely to be used to identify the natural person, account should be taken of all objective factors, such as the costs of and the amount of time required for identification, taking into consideration the available technology at the time of the processing and technological developments. The principles of data protection should therefore not apply to anonymous information, namely information which does not relate to an identified or identifiable natural person or to personal data rendered anonymous in such a manner that the data subject is not or no longer identifiable. This Regulation does not therefore concern the processing of such anonymous information, including for statistical or research purposes.

Official materials and guidelines

WP29, Opinion 4/2007 on the concept of personal data, 2007 (link)

WP29, Opinion 05/2014 on Anonymisation Techniques, 2014 (link)

EDPB, Guidelines 01/2025 on Pseudonymisation, 2025 (link)

Case	Data category	Personal data?
– 101/01 C Lindqvist (2003)	Names, contact details, job information, hobbies, medical information (injured foot) about parish colleagues published on a personal website.	Yes – individuals were named and identifiable; job/hobbies related to social identity; health info also covered.
C-141/12 & C 372/12 YS and M & S (2014)	(1) Factual data about asylum applicants (name, DOB, nationality, etc.); (2) Legal analysis within the same administrative minute.	(1) Yes – clearly identifiable data. (2) Unclear - legal analysis isn't "about" the person in the same way; its purpose is legal reasoning, not describing the individual (?).
– 582/14 C Brever (2016)	Dynamic IP addresses logged by a website operator (German Federal Institutions).	Yes - because the website operator can legally obtain identifying data from ISPs under certain conditions (e.g., cyberattacks).

CJEU cases

Case	Data category	Personal data?
– 434/16 C Nowak (2017)	Written answers submitted by a candidate during a professional exam and examiner's comments.	Yes - answers reflect knowledge, purpose is to evaluate abilities, and outcome affects rights; examiner's comments also relate directly to the candidate.
– 25/17 C Jehovan todistajat (2018)	Data from door-to-door preaching (names, addresses, religious beliefs, family circumstances).	Yes – names/addresses identify individuals: beliefs and circumstances relate to them and fall under data protection rules.
– 319/22 C Gesamtverband Autoteile Handel (Scania) (2023)	Vehicle Identification Numbers (VINs).	Yes - not personal data in itself, but becomes so if there is a database in which that number is linked to a person

'Processing'

> *Anything and everything.*

So, the first element: for the GDPR to apply, personal data must be 'processed'.

'Processing' == just about anything you could imagine doing with data: collecting, reading, storing, hashing, analyzing, sharing, archiving, damn, even just *deleting* — **think CRUD on steroids.**

> **Examples**
> A log line flashes through RAM for three milliseconds? Processing.
> Your QA engineer downloads a CSV, realises it's the wrong file and drags it straight to Trash? Processing. Technically twice.

Intent, duration, or depth of human involvement **do not matter.**

Virtually a permanent 'return true'.

'Involved'

> *Decision == control == responsibility.*

'Involved' is a phrase I made up to simplify things a bit for now, so just remember this: **for the GDPR to apply, the 'processing' does not need to be done by you.** As a matter of fact, you could be **held responsible to the data you have never seen and have no access to whatsoever.**

So **what** makes you responsible? **Your decisions/control over the process.**

> **Why?**
> Say, I don't take pictures of you, but I hire a spy to take them.
> Even if I then tell them to burn the photos, and have never seen them, it kind of makes sense I'd be responsible.
> My decision kicks it all off.
> My power, my control, my responsibility.

For now: processing == anything done with data. Your *decision* to use data == you are in control == your responsibility starts.

You are literally called a *'data controller'.*

Now, the courts interpret this insanely broadly.

In one case, a local library had a fan page on Facebook; they had access to an admin dashboard which let them see aggregated statistics, but no actual data on the individuals. A sane person would say that they are not, well, responsible for whatever Facebook does to make the stats. *Nay*, museth the Court of Justice, the fact that the library opened the Facebook page meant that those searching for it, even without a Facebook account, would perhaps click on a link to the page, and thus, the library would facilitate Facebook's processing of data. Meaning — joint liability.

But don't get too caught up in this, because, frankly, it makes no sense.

Start with the basics: you have at some point decided *something* about this data, or your actions somehow play a role in it being processed? return true.

(There is another scenario – you *do* happen to store some data, but you don't get to decide anything about it: you are just following someone's instructions. Say, you make an app for a business client which then tells its employee to use it to punch in/out. They can ask you to delete all the data. You have no control over how it is used (say, to calculate salaries). In this case, you do not have the *same* level of responsibility - you are considered a 'data processor', because you do not get to call the shots. Lots of stuff in the rest of the guide would not apply to you when it comes to those activities — where you just do stuff for someone else (no, the end-users don't count). But that's a bit complicated, so let's put a pin in it for now. Stick to decision == responsibility for now. We'll touch upon data processors a bit later. And, for the record: no, 'but a user can tell me to delete the data, so I am not a controller' does not fly.)

'Personal'?

> *Everything.*

GDPR applies to processing of 'personal data'.

Now, this is where people trip up. **This is IMPORTANT!**

'Personal data' isn't (only) stuff like names or e-mails. It's **anything** that lets you, or **anyone** else in the world, in **any** remotely possible way (including with combining with other data), **pick a real person out of the crowd.**

> User ID like "f390ad0f4"? Personal data.
> Dynamic IP addresses? Yup, according to the EU's Court of Justice, even if you have no idea who's behind it. Somebody (usually, the ISP) does.
> Device IDs, order numbers, location pings? You've guessed it.

Again, the question is **not** what kind of data it is, or whether you know whom it belongs to. **Can someone, anyone, anyhow, with any additional amount of data, figure out who this is, now or at some later point? Personal.** Meaning — yes, the GDPR applies.

Doesn't matter if it's medical records or a freaking session id.

It is *idiotic*.

But it is what it is.

So... return true.

For the nerds: Yes, there is an exception, and yes, there is the recent SSB case. My take: whenever in doubt, assume that the GDPR applies. SSB seems to be extremely strictly interpreted by the EDPB. Let's see when the Digital Fairness Act comes.

Anonymization

> *It's a lie. Virtually.*

And, as a side note: when people say 'anonymized dataset', there is an extremely high likelihood that they are **wrong**. In practice, most organizations don't actually anonymize-they '**pseudonymize**' - i.e. make it **more difficult** to identify someone.

In other words, the GDPR likely applies to the vast majority of what you, and any reasonable person, might call 'anonymous'.

> Don't believe me?
> In 2006, Netflix put out an anonymized dataset with over 100 million movie ratings from about 500,000 users—each entry had a user ID (anonymized), movie ID, rating (1–5 stars), and the date it was rated. Sounds anonymous?
> **Nope.** Researchers found that by comparing just a few ratings and dates with public user profiles on IMDb, they could figure out who the users were. So be very, very skeptical of the word 'anonymization'.

The good news is...

This is where the bad news end.

Even though the GDPR applies to virtually any piece of information, it's not impossible to comply with it (and for the most part, not super *difficult*, either).

So let's roll on!

Recap time!

The GDPR does not stop to apply because you don't know who the user is.

It applies to *any* data which can in *any* way be linked to a person, by *anyone*.

Encrypted, minimal, meaningless?

Personal.

Not even storing/sending/accessing it? Doesn't matter.

I know.

```
  @Observable
final class User: Identifiable {
    var name: String  // Personal data
    var sessionId: UUID      // !Also personal data!
    init(name: String, sessionId: UUID = .init()) {
        self.name = name
        self.sessionId = sessionId
    }
}
```

02

What do I have to do?

02.1

Know your purpose

Where to start?

> *With the basics.*

Here come some good news: while the GDPR is so long that I *absolutely* suspect the drafters were getting paid by the sentence, we have a **single** article – **Article 5** – which **contains all the ground rules.**

We call them GDPR '**principles**', and complying with them is...

Well, practically 90% of GDPR compliance.

Better yet: those rules are, for the most part, honestly just **common sense.**

Think of them like using `import Foundation` and your primitive types. Without them, everything falls apart. But understanding what they do isn't hard. And there is no need to invent a new type where a simple boolean would do.

... Aaaand, yes, there is a catch.

> *Because of course there is.*

The principles are **awesome**. And **simple**.

But before we can meaningfully look at them, you *will have to* do some work.

If you think of the GDPR principles as functions in code, they all require **the same input parameter**: *the reason(s) why* you are processing people's data (also called the 'purpose(s) of processing').

I cannot stress this enough. So let me repeat it:

> ⚠ Before you do ***anything*** under the GDPR you need **a list of goals** for which you use the data. Each one of those is called a 'purpose'.
> *'What am I doing? Why?'* are the questions you need to ask. And, well, answer.

We will work on it in the next part.

For now... Think of the GDPR as the world's most tightly coupled codebase. One thing falls apart, it all falls apart. (And force unwrapping is about as good idea as always.)

```
func doAllTheGDPRStuff(for data: PersonalData, purposes: [Purpose]?) {

    guard let purposes = purposes, !purposes.isEmpty else {
        fatalError("No purposes, can't invoke other functions!")
    }

    for purpose in purposes {
        minimizeData(for: data, for: purpose)
        defineRetentionPeriod(for: data, for: purpose)
        secureData(for: data, for: purpose)
        //...
    }

}
```

The GDPR's MOST important step

> Is to know *why* you are using data.
> This is your 'purpose'. And it shapes *everything.*

Do you have to ask for consent? Which rights do people have? Which data can you gather? Can you re-use or share it? When do you need to delete? What kind of security do you need?

It depends.

On what?

On **why** you are using that data.

You need **specific**, **concrete** reasons why you need the data — or goals you are trying to achieve – stated in a sentence an average grandma can understand.

Like: "I need data so that I can _____, _____ and _____."

How do I start?
Think... Features.

For example: "*I am making an app to let people read CJEU judgements. I will also let them make highlights, bookmark things, and sync across devices. I will also need to do some debugging.*"

Honestly? Go through the application lifecycle and jot down *why* you are using data.

Even if it's so simple as *"I let people type in their search query in my SearchView so that they can... search".*

Homework time!

> Please stop reading for a second. Or like 480 seconds. *Please.*

YOUR JOB:

Make a **list of reasons** *why* you are using people's data.

It will be *impossible* to continue without it.

Yes, I know it's boring. But it *has* to be done.

A cup of coffee, a deep breath, and a few existential questions on the meaning of it all, and we are good to go!

How specific should my purposes be?

Very specific, but not absurdly so.

"Developing new products and services" is too vague.

"To send e-mails on new app features" is good.

"Sending e-mail #472-B with A/B testing variant 3 on subject line effectiveness regarding our new feature X" is excessive.

Small advice

DO NOT get bogged down into data categories just yet.

It's tempting to jump straight to 'I am using user e-mails to...', but that's just gonna frustrate you. Start from the big picture – goals – and then we will drill our way down to data categories later. Don't fuss over variables while you still don't know if you even need a class/struct to begin with.

In other words, ignore *which* data you are using. Just document *why* you need to use the data.

Do you have the list? *I am watching you.*[1]

Yaaaay! Quick recap!!

You are *awesome!!!*

Quick recap update and let's move on:

> **Chapter 1:** GDPR applies if you are established in the EU or trying to sell stuff/monitoring people in the EU. It applies whenever you decide people's data will be used in some way, or store/process it for others. 'Data' means anything, like literally anything, even a damn UUID which exists for 3ms.

> **Chapter 2:** To know what you need to do under the GDPR, you must start with a list of specific reasons why you are using people's data. Each one is called a 'purpose'.

Now, let's dive into the *actual* requirements!

(He said, 30 pages in.)

02.2

A matter of principle

Let's get to the good stuff!

By 'good stuff', I mean the principles I mentioned earlier.

Remember? When I said that they are **practically 90% of GDPR compliance?**

Well, coming up next!

We're going to learn about the principles, then take our **list of purposes which you absolutely have** and work through them individually. The idea is straightforward — for each purpose we mentioned, we do a check against the principles.

Spoiler alert. There are **six** principles in total.

In fancy terms, they are called 'lawfulness, fairness and transparency', 'purpose limitation', 'data minimization', 'integrity, confidentiality and security', 'accuracy', 'accountability'.

We will try to give them some more reasonable names.

Even 'Rumpelstiltskin' is a win at this stage, but hey.

'Purpose limitation'

> Data without any purpose == violation.

Now, here is the first 'real' rule: **you have some data laying around, but have no purpose for it?** GDPR violation. **Delete delete delete.**

Shipping an app with dead code is bad practice. Just extend that logic to the GDPR: processing people's data '*well... because... I...*' is illegal.

Yes, that might sound like the GDPR is way too strict.

But think about it: why have things you don't need?

You can show to your boss you actually need 39 external libraries? Sure, import away! But don't go about importing an XML parser into an app which only deals with JSON. It serves *no purpose*.

You know, let's give it a positive spin!

What can you collect? Anything, if you have a (legal) purpose. Other requirements *will* follow, for sure, and some might be a bit tough, but that's our *starting* principle — and it's a good one!

So... delete the stuff which *you have no reason to have* as soon as you get a chance. Let's remember this one simply as 'Know Thy Purpose'.

Ah, the sheer cosmic depth of the Regulation 2016/679.

> **Example**
> For my app for highlighting EU judgements, I do not need to ask people about their horoscope signs – as the data serves no purpose whatsoever. However, if I were to introduce such a purpose ("find a CJEU ruling matching your lucky stars"), it'd be fine to collect it, assuming everything else is taken care of.

▼ **Changing the purpose along the way?**

> Now, some of you might say, what if I collected data for one purpose, but want to use it for something else?
> Say, you have actual user logs, and you want to use them to test your database schema migration?
> Well, principally, you should not (for reasons both legal and technical) — but there is no absolute prohibition against re-purposing.
> Re-purposing is **allowed**, in other words, as long as you can show that the **new purpose is *compatible*** with the old one — based on their goals and user's expectations. It's too much to get into it here, but basically, think: did I get their e-mail address to give them customer support, but now want to send marketing e-mails? Incompatible; I need new consent/new legal basis. My boss gathered ID card logs to secure the building, but now wants to use them to check how often I am at work? Incompatible.
> Rule of thumb: if you'd find it, as a user, surprising that the data is being used re-used to do *abcd*, don't do *abcd*.

'Data minimization'

> *Don't be, like, crazy.*

This one is nearly identical, but with a slightly different spin.

It basically says – if you **can** achieve your purpose **without** some specific data, **do it without** that data.

> **Why a separate principle?**
> Think of it like this.
> I want to serve you a dinner, and I want to know about your allergies. Protecting your health == clear purpose == I am fine on the purpose limitation principle we've just seen.
> *But*... There are many ways I could achieve that purpose.
> I could ask for your full medical history. I could review the footage of you at a grocery store and see what kind of food you avoid. I could, if feeling particularly murderous, put some peanut butter into your meal and ever-so-carefully film what happens next.
> Sure, I would achieve my purpose (at least in the first two cases), but in a way so unnecessarily intrusive that it makes sense to no-one.
> Hence, illegal. Hence, delete.

Code equivalent? You include TensorFlow to detect if a string contains "hello". Yes, it detects it. No, your laptop's fan will never turn off again.

Ask yourself: can this exact thing be done with less data? If yes, delete what you don't need.

'Storage limitation'

> *You no longer need it == delete it.*

When should you delete the data? **When you no longer need it.**

There is **no such thing** as specific deadlines under the GDPR – 90 days, 2 months, 3 years. As long as you need something, keep it.

... 'But what if it might be useful someday?', I hear you saying. Well, hate to break it to you, but that fits the American Psychiatric Association's definition of hoarding to the letter. Don't need it? Get rid of it. Like, yesterday.

That being said...

Don't get too hasty when setting up deletion rules. Why? Think... SwiftUI. You notice your view model has a function which is not, well, used in a particular view. You delete the function. Have you broken 10 other things? So the same way you would click 'show callers' on a function, when setting up a rule to delete personal data, make sure it is **truly** not needed for **any** purpose at all. Otherwise, the mess ensues.

> **Examples**
> I am blanking on examples, but, say, I am grading exams. (Funny that comes to mind.) I enter the grades into our proprietary database. (May it burn in Hell.) Then, 'because the GDPR says so', someone sets up a strict deletion rule – the grades are vaporized three months after the student's deadline to request a new assessment expires.
> But, wait. Surely we need them to... I don't know, correct errors/grade entry mistakes? Check for plagiarism if someone reports them in 3 years?

So the rule is simple: the GDPR will **never force you to delete the data when you can show that you need it**. When designing a delete rule, just ask - how long will I need it?

Now, a small break!

Get some candy (or an apple, if you are so inclined) — and we're gonna cover a few more principles, then dive into some examples.

Coming up next, *security*.

'Integrity & confidentiality'

> *Lock your doors at night, but don't build a nuclear shelter.*
> *(Unless a normal person would, too. Which ummm these days...)*

Great! So you've embraced your inner minimalist, and have gotten rid of a bunch of data. What's next?

Well, you have to protect the stuff you kept.

The GDPR calls this the principle of 'integrity and confidentiality'. Normal people call it 'security'.

Now, don't freak out. Don't rush off to buy that OurBaseBase Pro Premium Plus XXL storage plan with 'guaranteed compliance' and 'real-time monitoring'

Because the GDPR **doesn't care** if you use AWS, a server in your basement, or a carrier pigeon — as long as the setup is ***reasonably secure*** for what you are doing, in the ***grand scheme of things***.

▼ **A small rant**

> About those"GDPR-certified" platforms...
> DO NOT. *EVER*. BUY. THOSE. THINGS. (For the sole reason of GDPR compliance, that is.)
> There is no such thing as a "GDPR-compliance-fix-tool".
> Anyone who tries to sell you one is selling snake oil, and you should run in the other direction. (Or, depending on your personal inclination, run them over. I am leaning towards the latter, as what they offer is pure fraud.[1])

Let me repeat that: the GDPR does *not* require some crazy security standard, and most *certainly* does not require a specific technological solution across the board.

Instead, it looks into the ***context:*** what kind of data are you using, in what kind of an app, for which purpose? And simply says: **do what is reasonable**.

Don't believe me? Here:

▼ **Article 5**

> Personal data shall be [...] processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures ('integrity and confidentiality').

Or, after a quick refactor:

You should process data in a way that ensures *appropriate* security — and you need to implement *tech* solutions (e.g. software) and *organizational* procedures (e.g. employee policies) to the degree that makes sense to a reasonable person.

How do I know if a measure is reasonable? *Plain common sense.*

Example
Recipe-reading app? Well, make sure not to store usernames in plaintext.
An app letting a doctor do a video consultation? Biometrics, quadruple authentication, mandatory user mugshots every other month.
OK, not that far, but you get the idea.

Let's look into the details.

Measuring the risk

> *Yours is bigger.*

For a long time, my approach to risk-measuring was the basic 'what's the worst that could happen?'-type of brainstorming. But I have, after many a meeting with people with thin watches and briefcases (henceforth known as 'lawyers'), and those with ~~redacted~~ and ~~redacted~~ (also known as 'CISOs'), come to conclude that it does *not* work.

Why?

Because you are likely either going to overshoot (I kid you not, one person, under risks, listed 'global warming'),[1] or undershoot ('which hacker is gonna target me?', said a certain grave).[2]

Two things instead.

Start with some empathy. **How much would *you*, as a user, care if this data got out?** It's not representative of the risk — you may be making an app for the elderly while still at high school — but it sets you in the right mindset. Then, research time: start off by looking into whether there have been any reported attacks against **services comparable to yours** — again, it's not a full picture, but it gives you a *start*. If you jot those two down in broad terms — you have impact times likelihood, which, people tell me, is a risk assessment. A *crude* one, but one *nonetheless*.

> **Super, super, super important point:** when assessing the risk, you must do it **from the perspective of a person** whose data you are processing, and **not**, I repeat, **not**, from your own perspective. You can't just say *'yeah, look, I am ISO-certified'*. Fine as that might be for your company (although, realistically, meh), it **does not** give you a free pass from the GDPR. Perhaps the measures you have implemented under it are sufficient — *perhaps* — but you cannot possibly know until you have evaluated things from your user's POV.

Mitigations

> *Try your best.*

Once you have identified the risks, you need *specific* measures to address it — as mentioned, technological or organizational.

Tech-wise, I am certainly not the one to be telling *you* what works. An obvious candidate is strong and reliable encryption (http:// is an exception, not a rule in your app, hopefully). Where possible, prefer local storage (assuming function parity). You likely do both anyway. Check for bugs/common points of data exfiltration (yeah, I once hardcoded an OpenAI API key on a client — *don't* be me).
Policy-wise: teach your colleagues that ^⌘Q (or Win+L) are a must... by setting embarrassing wallpapers on any device you find unlocked. Send them a *phishing* email stating 'you have been phished' (works on those who normally know better than to claim their inheritance from a prince). You work alone? If coding outside your home/office, turn FileVault on. If you print sensitive documents, shred them once you are done.

> **When is a measure appropriate? Is it too little, too much?**
> It depends.
> Don't hate me, because it *literally does.*
> Remember, we legally look at the state of the art (do people still use Caesar cypher?); the costs of security measures (even the EU doesn't require you to buy a quantum computer to protect a beer pong app); how much data we process (someone's IP address or their entire cloud storage); our purposes (am I storing their phone number to call them in case of a medical emergency); and importantly, the risks that my processing can pose for people.

You *probably* do most of this stuff anyway. But things can always be done better, and hey, you would *not* be spending your time reading a damn GDPR guide from a random guy online if you *didn't* care, so — pen, paper, a quick sanity check?

Speaking of pen and paper...

'Accountability'

> *Jot it down!*

There is *One Principle* which kinda wants to rule them all: 'accountability'. It just means one thing – you have done everything we talked about so far?

You have to prove it.

The legal logic being... **You need to be able to show that you are compliant, and if you can't, you are not compliant.**

So open up your text editor. Write 3 sentences. Come back tomorrow.

Remember: it does not need to be a legal document. **It does not need to be 200 pages.** Things you might have jotted down while reading this? *Potentially enough.*

▼ **How in the world is this relevant in practice?**

> Oh, but it is.
>
> Data protection authorities might not have the time, resources, or, for that matter, the expertise to audit your code security, but they sure as hell can prove you did not document it, making it a super low-hanging fruit for them. (Yup, this was how Facebook ended up being fined for improper security measures. They could not prove the codebase was not secure. But they could prove lack of docs.)

Good codebase == good docs. Same here. Good GDPR == good docs. Where reasonable, a one-pager in your editor of choice. Where reasonable, 80 pages. But have *something*. At least 'Here are my purposes, this is when I delete the data, this is how i secure it'.

ALMOST DONE.

'Lawfulness, fairness, transparency'

> Don't be an *******.

Last one! (We're on principles, remember? Have a purpose, don't be a Smaug, implement normal-person-security?)

So... Lawfulness, fairness, transparency.

OK, I am not gonna take '*lawfulness*' here because it has a special meaning (think 'consent' et al), so let's do 'fairness' and 'transparency'.

'Fairness' == don't be evil. (Aged like milk.)

> Dark patterns in UI design? Unfair. Using data to figure out what time of the day people drink, so that you can push them a subscription pop-up? Unfair. Asking for consent five billion times? Unfair, even if they say 'yes'. Follow your gut feelings, and, like, don't be an ass.

'Transparency'? **Let people know what you are actually doing.**

Now, there are several components to this – and we will look at them a bit more when talking about privacy policies – but basically, **talk like a normal human being** and don't go around plotting clandestine schemes.

And before you say 'no, I have to use difficult *legal* language in my privacy policy because of *liAbIlIty*', whomever told you that is an idiot.

So: simple, clear info! We're gonna practice soon.

But hey? You are almost done with Article 5!

See, not *horrible!*

But yeah, let's recap.

90% of compliance is...

Step 1:

List all the reasons why you are using data

To let people highlight legal texts, cross-sync them, pay for a 'pro' version, receive marketing e-mails.

Step 2:

Any data you have without reason -> delete.

User's gender, location, favorite color seem pretty useless to me. Off it goes. I'll get it when/if I end up needing it.

Step 3:

Any data you technically have a reason for, but you might just as well have done without it -> delete

Mother's maiden name for security? Surely there is a way I can secure stuff without it?

Step 4:

Any data-related security risks -> mitigate (reasonably)

User highlights in public judgements? Yeah, maybe don't use password 'test' but don't lose sleep over it. Lawyer's notes typed about the same judgement? Lockdown.

Step 5:

Shady stuff -> remove; doing something -> inform

A fake '20% discount' if they consent to my marketing? Yeah, no. Started using live data for testing for v. 2.0? Yup, inform.

Step 6:

Document steps 1-5

For my app, this would pretty much do. It's a safe bet yours is more complex. But the logic is the same.

```swift
// Basically

struct QuizAnswerLog {
    let userID: UUID           // need it to link to user, ok
    let age: Int               // why? don't seem to need it anywhere rn
    let deviceFingerprint: String   // hm, overkill?
    let questionID: UUID       // need it for linking, ok
    let answerText: String     // what the user typed in, need it to log the answer, ok
    let timeTaken: Double      // time user took to answer - delete if never actually used
    let debugToken: String?    // delete when done debugging/after reasonable time w/o bugs
}

// Not a quiz app?? How did it get here?? Delete the whole damn thing
```

Bonus points? Consistency.

There! That was the toughest part!

Are you “GDPR compliant”? Not *just* yet (partially because no such thing exists, and yes, I shall die on that hill), but you are A LONG way down the path.

So the whole thing so far is “*delete the stuff I don’t need, document what I do, be reasonable about security, don’t be shady*”.

Which is not all that hard.

The thing is... It's like going to the gym. In January, I am there. In February, mwhahah LOL *no*.

What you really need is a plan to do the GDPR stuff *periodically* and *consistently*.

> **Yeah, but how?**
> Put a date in your calendar to run a sanity check on which data you have and why.
> Adding a new API? Check its security and take a quick note on it. Even if it's just 'looks decent by the modern standards'.
> Removing a feature? Remove the associated data. If sure you don't need it anymore.
> It’s just like building good coding habits – you need something resembling a process. You need an equivalent of ‘every time I refactor, I commit’.
> Again, even if just to write ‘I checked, all is well’.

Once you do that?

Pop *some champagne!*

In lawyer speak, you have “*ensured compliance with purpose limitation and data minimization principles, established data retention periods, implemented appropriate technological and organizational measures to mitigate security risks to fundamental rights and freedoms, and demonstrated a high level of accountability*”*!*

:the EU nods approvingly at you:
:while inventing 18 other regulations which are far, far worse:

Let's update our master recap!

> **Chapter 1:** GDPR applies if you are established in the EU or trying to sell stuff/monitoring people in the EU. It applies whenever you decide people's data will be used in some way, or store/process it for others. 'Data' means anything, like literally anything, even a damn session id which exists for 3ms.

> **Chapter 2:** To know what you need to do under the GDPR, you must start with a list of specific reasons why you are using people's data. Each one is called a 'purpose'.
>
> To comply with the core rules: first, delete all the stuff you don't truly need, and make sure you set up a reasonable deletion routine per each purpose. Secondly, secure the data depending on how much of a risk it poses, don't be an ass, and give people proper info. Thirdly, be careful if re-using data collected for one purpose for something unrelated. Not necessarily in this order.

TBC.

Break!

But look, if you have been reading so far in one go, take a walk, a break, a nap.

What remains is not *that* tough, but if you have just crammed everything from extraterritoriality and anonymization to data controllership and purpose limitation into your head, your brain is very, very likely overflowing. You deserve some rest.

But *do* come back later!

03

Legal basis

Legal basis

> *You need to draw one of six straws for each purpose.*

So!

You've got your list of purposes (**good!**), you've minimized your data (**great!)**, you've set up deletion schedules (**excellent!!**).

But there's one more thing the GDPR wants to know: what gives you the right to process this data in the first place?

This is 'lawfulness' – the principle I shamefully (yet strategically) skipped over earlier.

It works like this: GDPR gives you a list of options ('legal basis'), and you must pick one for *each* purpose. Almost like a multiple choice question where you have to pick an answer for each reason why you are using data. You can't leave it blank. You can't pick "other" and write in your own.

You get exactly six options to choose from.

And yes, one of those is consent, but it is only one and I will rage about it soon.

For now... Think of it like an API that only accepts specific enum values. You can't just send random strings and hope it works. Article 6 gives you six valid legal bases - and you need to map each of your purposes to exactly one of them.

The rule is simple: 1 purpose = 1 legal basis. **You take the list of purposes you have made, and you add a legal basis next to it.**

Let's see what the options are!

Pick your poison

> Ready? Here are your choices.

Here are your picks. Don't get intimidated. **We will look into them in detail,** but basically, think: which one would make most sense for each of your purposes? You use data to... Let users sign up, monitor for security breaches, profile with ads. For each one of those, pick one ground. **Can't find one? Delete.**

1	**Consent** They explicitly said "yes, use my data to... *abcd.*"
2	**Contract** You need their data to deliver t-shirts they purchased? You need to make it possible to submit a quiz answer in a quiz app? Fair game – if you don't use data, you'd be breaching a contract, so no need for consent. Think like this: my bank does not need my consent to send me a credit card bill.
3	**Legal obligation** The government says you must. Tax records, anti-money laundering stuff. You don't get a choice; neither do they, so you can process away — no consent, no contracts!
4	**Vital interests** Life or death. Literally. Medical emergencies. Don't get creative with this one.
5	**Public task** You're a government doing government things. Probably not vou.
6	**Legitimate interests** The Swiss Army knife of legal bases. You have a good reason, they'd reasonably expect it, and it's not creepy or risky. Analytics? Fraud prevention? Security logs? Usually fits here. The gist of it is – no need to pester for consent if what you are doing is not gonna harm users more than it helps you, and a reasonable, decent person wouldn't find it surprisingly shocking. Super flexible, but demands *real* docs and transparency.

03.1

Consent

GDPR != consent

> *... And that is good.*

So you've decided you need consent for something. Marketing emails, maybe. Or app analytics that go beyond the basics.

Good for you.

But here's the thing: consent isn't just a checkbox.

A user said 'yes'? Hate to break it to you: that does not a valid consent make, legal Padawan.

Consent has got four very specific requirements, and messing up any one of them makes the whole thing invalid. It's like a function that returns false if any parameter is false or nil. 'Yes' does not mean 'yes'; not under the GDPR — unless all the requirements fall into place.

That is, partially, why it sucks, but I'll get back to you on that one.

Still, rule of thumb: if you *can* avoid asking for consent, *do* avoid asking for consent,[1] because, to quote the eternal wisdom of Veep, it does not work, and it makes a mess.

Saying 'yes' is not enough

> *Not nearly, in fact.*

To re-iterate: just because somebody tapped on 'yes' does not mean you have obtained a valid consent. For consent to be a valid basis it *must* meet certain legal criteria.

Freely given.

Specific.

Informed.

Unambiguous.

We are gonna disentangle them slowly. Ye shall wake up at night screaming of Article 4, or, on a bad one, 7.

Nah, not really.

For now, just remember: when you decide to go for consent, simply creating a fancy UI and logging yes/no does not fully cut it – we have to — you guessed it — unpack it one thing at a time.

So let's go!

'Freely given'

> Actual, meaningful choice.

"Freely given" means people *actually* have a choice. Not a theoretical choice. Not a "well, technically they could refuse" choice. A real, practical choice **without *any* negative consequence for saying no *whatsoever*.**

> Bad: "Accept all cookies or you can't use our site"
> Bad: "Sign up for our newsletter to create an account"
> Bad: "We'll only give you customer support if you consent to marketing"
> Bad: "All your posts will be public unless you pay."

Think of it like offering someone cake. If you say *"eat this cake or I'll fire you"*, that's not freely given consent to eating cake. That's coercion. Even if they would, in theory, like cake.

So, the real test is: **can someone refuse consent and still get the core service you're offering?** If not, your consent probably isn't freely given.

Also, and this is easy to miss, **no bundling** of consent. You can't ask for permission to send newsletters AND to track users for analytics in one checkbox. That's like asking "Do you consent to A, B, C, and D?" — how do they opt out of just C? It also breaches our next principle.

Also also, to assess whether consent is 'freely given' you need to look at **the entire context** – like the relationship between the parties. Will they be afraid to say 'no'? *Unlikely* when it comes to an app, but, if your boss ever asks you to consent to some use of your data – their processing might be illegal, even if you say 'yes': very few people would feel free to say no to their boss without fearing any consequences. Just one of those 'oh, nice to know' things to file in your mental cabinet for the day you decide you wanna be difficult.

'Specific' and 'informed'

> *One purpose, one consent. Normal language.*

Let's do 'specific' first.

Each consent request should cover exactly one specific purpose.

> Bad: "I consent to data processing for business purposes"
> Bad: "I agree to receive communications from Example Corp"
> Bad: "I consent to cookies and similar technologies"
> Good: "I consent to receiving weekly newsletters about new product features"

Be as specific as you'd want a junior developer to be when describing what their function does. "This function processes data" doesn't tell you much. "This function sends welcome emails to new users" — now we're talking.

As for 'informed', we'll talk more when we get to privacy policies, but people **need to know what they're consenting to** before they consent. Shocking concept, I know.

> Bad: Checkbox next to "I agree to the Terms and Privacy Policy" (20-page linked documents) Bad: "We use cookies to improve your experience" (vague as Hell)
> Good: "We'll use your email to send you our weekly newsletter. You can unsubscribe anytime. We'll keep your email until you unsubscribe or after 2 years of inactivity."

Make it like good code comments — enough detail to understand what's happening without having to read the entire implementation.

As you can see, these two already correspond to what we've learned about the principles. ('Be reasonable.')

'Unambiguous'

> *Basically, active.*

This is EU-speak for "not a dark pattern." People need to actively and clearly indicate they consent. No pre-ticked boxes. No "continuing to use our site means you agree." No double negatives like "uncheck this box if you DON'T want marketing."

Bad: Pre-ticked "I accept selling my soul"-box
Bad: "By continuing, you consent to..."
Bad: "Uncheck this to opt out"
Good: Empty checkbox they need to actively tick
Good: "Yes, send me newsletters" button

Think of it like confirming a destructive action in your app. You wouldn't pre-select "Yes, delete all my data" - you make users explicitly choose it. It just requires that a user does something.

So what?

> *Check if you need consent.*

I know I have stated this very clearly before, precisely 87 times, but I need to repeat it. **Stop defaulting to consent.**

Remember how I said there are good alternatives? We will talk about those next.

Consent is the diva of legal bases — high maintenance, can disappear overnight, and everyone thinks they need it when they probably don't.

Sure, it feels fancy and "user-empowering," but now you're stuck managing consent withdrawals, re-consent flows, and that passive-aggressive GDPR banner that covers half your homepage.

"But users like having control!"

Sure. They also like not getting pop-ups every five seconds asking permission to breathe.

Let's just do a sanity check: not only does consent have those sky-high requirements you've just seen, but remember, it can always be withdrawn. Which creates a logistical nightmare.[1]

What do you do if you *need* this data for a vital functionality, but they stop consenting?

Bring out the alternatives!

:fanfare:

But first! Deep breath!

> *Time to be proud!*

So by now, you know tons more about GDPR consent than your average lawyer! *Seriously!*

You are one of the very few people who know that tapping 'I agree' on a pop-up 'I consent to use of my data to improve our services and develop new ones' is pure nonsense – even if they tap 'yes', consent is totally bogus. Useless. Nil.

> ***Quick test!*** **Why?**
>
> It's neither freely given nor specific (1 consent for 2 purposes), and, for that matter, kinda not informed either (WTF does this actually entail?)

And if you have noticed another issue – that this kind of purpose talk is gonna get you in a world of hurt under Article 5 (principles) – how do you define a deletion routine, secure or minimize data for 'developing new services'? — then you are a GDPR rockstar already. Truly!!

And WOW, if you have noticed *another* thing – that asking someone to tap on 'I consent to Privacy Policy' is pure useless nonsense, for the reasons mentioned above (a person consents to a *purpose*, and not to a freaking policy) – you are my favorite person in the world. And you should seriously consider studying to be a lawyer, if you are not one. I am not joking.

Your favorite time! Recap update time!

> **Chapter 1:** GDPR applies if you are established in the EU or trying to sell stuff/monitoring people in the EU. It applies whenever you decide people's data will be used in some way, or store/process it for others. 'Data' means anything, like literally anything, even a damn session id which exists for 3ms.

> **Chapter 2:** To know what you need to do under the GDPR, you must start with a list of specific reasons why you are using people's data. Each one is called a 'purpose'.
>
> To comply with the core rules: first, delete all the stuff you don't truly need, and make sure you set up a reasonable deletion routine per each purpose. Secondly, secure the data depending on how much of a risk it poses, don't be an ass, and give people proper info. Thirdly, be careful if re-using data collected for one purpose for something unrelated. Not necessarily in this order.

> **Chapter 3:** For each purpose you have, you need a legal basis – one of the six alternatives. One 'to let users log on', one to 'send marketing e-mails', and so on.
>
> Consent is an alternative most people talk about. But consent is tricky. *Nothing* negative can happen to a user if they say no (disabled feature), it can be withdrawn whenever, it must be granular, informed, active. Avoid when in doubt.

TBC.

03.2

Alternatives to consent

Contractual necessity

> *You can use the data to deliver what you promised.*

When you sell an app, or when people visit your website, you enter into a contract. To fulfill your duties under that contract, you must use some data. This means that **for the majority of core functions in your app, there is no need to bother with consent or other legal grounds.**

The app would truly break without this data? You are good to go.

> Picture this: you are selling a subscription to an app which lets users track how often they exercise. Do you need their consent to process this data? No. Why? Because it would be *impossible* for you to deliver the service without such data: if a user was to 'decline' giving their consent, your app could not deliver its very core feature.
> It's easier if you think of it as ordering food online: a restaurant must get your address, or they would not be able to deliver the food you have ordered - therefore, no consent is needed.

However, **do be careful**. The test for using this legal ground has become very restrictive: it only applies to core features, and only if the data is 'indispensable' - that is, a 'must-have' rather than 'nice-to-have'.

That approach in itself is pretty idiotic, for reasons I am not going to get into here, but if you meet me at a party, I will tell you all about it. As others look in horror.

Anyway. Let's think practically.

> What *can't* your core app function without?
> I am thinking... User logins/e-mails/passwords. Payment information. Info like 'user A shared this doc with user B and C' – as to make collaboration possible. Photos – to make it possible to edit them in a photo-editing app. Preferred wake-up time – to make it possible to use a custom alarm/scheduling app. You get the gist.

So, if you see you are relying on consent for a must-have feature, drop it – go with contractual necessity.

BTW, pro tip: remember how you must delete data when someone withdraws their consent? If your legal basis is contractual necessity, users have no right to ask for data to be deleted while the contract is in place. (Makes sense –I can't ask the restaurant to delete my address before the food is delivered. Or my bank to delete all my info before I pay my credit card bills.) So yup, it solves a few deletion problems for you as well.

Legitimate interest

> *Please remember this one.*

You can use data without consent **when it makes sense and hurts no-one.**

Legitimate interest means you are using for a reasonable, fair reason; that people can reasonably expect it; and that hurts them less than it helps you.

That is the elevator pitch. The GDPR itself hides the same idea behind forty-seven commas in Article 6(1)(f) .

I have spared you the commas.

Now, why is this important?

Because this is the GDPR's way of being sensible.

> Say there is a local library, and say they want to see how many people borrow books on the GDPR, and how long they keep them. They do this so that they can optimize future purchases, decide what kind of seminars to organize, whatever. Now, can they call each person who has borrowed a book and ask them 'hey, do you consent?' – I mean, sure. But that would be stupid, so, even the GDPR says, 'yeah, this one is fine, just do it! – so, as you can imagine, the practical implications are huge.

But of course, there's a test.

Actually *three* tests.

Legitimate interest tests

Test 1: Do you have an actual, meaningful, specific reason?

You need to show which interest of yours — or others — using the data serves.

Good reasons:
Security monitoring (someone's trying to brute-force your API)
Crash reporting (your app shouldn't randomly explode)
Fraud detection (that's 47th login attempts from North Korea in 3 minutes)
Internal debugging (why did the payment flow break?)

In principle, the threshold is rather low — making money *is* a legitimate interest, but you do have to be able to articulate it. "I need to do this because..."

Test 2: Do you really need the data?

You need to show that the data you are gathering is, in fact, necessary to achieve that interest.

Not in a 'it will be *impossible* for me to do any debugging if I don't get this data'-way.

More like 'this is a sensible way to do it, and I can show that this data really, really helps me'-kinda thing.

Test 3: Is people's privacy more important than your interest? Do they expect stuff like this to happen?

This one's pretty intuitive.

If your processing would genuinely freak out a reasonable person, you're probably doing it wrong. First, let people know what you are doing with their data (you have to do it regardless of legitimate interest or not). Secondly, think — should I *really* be recording every. single. possible. thing. a colleague does on their PC to make sure they don't exfiltrate documents?

Well, obviously not. *Unless...* That person is already suspected of selling confidential info, and there is no other way you can find out if they are guilty, and you do it for as minimal amount of time as possible. *Or unless...* You work in a nuclear missile silo.

I hope you get the gist. Their interest, your interest, circumstances — pass as input. Weigh them against each other. How important is what you are trying to do compared to how invasive the processing is? In other words — it's a balancing test.

Just think with your common sense.
Preventing your servers from being hacked? High necessity, low creep factor.
Fixing critical bugs in your payment system? High necessity, obvious benefit to users.
Secretly recording everything users do "for quality purposes"? Low necessity, high creep factor.

More examples

> No consent. Legitimate interest. All the way.

Server monitoring – You log IP addresses, timestamps, and error codes to catch security issues and system failures. Classic legitimate interest. You need this stuff to keep things running, users expect you to keep things running, it has a minimal privacy impact. If you go wild with monitoring? Well, breach remember data minimization principle?

Crash reporting – App crashes, you automatically collect stack traces and device info to fix the bug. Pretty reasonable. Users hate crashes more than they hate you knowing their iPhone model.

API abuse detection - Someone's hammering your endpoints suspiciously. You flag the pattern and maybe rate-limit them. This falls under "keeping things working for everyone else."

Database backup integrity — You need to test restore procedures, which means processing live data in your staging environment. As long as you're not doing anything weird with it, this counts.

Internal debugging – Production issue happens, you need to trace through logs containing user IDs to figure out what broke. Temporary, purposeful, directly related to service delivery.

Remember the docs!

> *Pics or it didn't happen.*

Want to use legitimate interest? Awesome!

Just please remember to document:

> What you're doing;
> Why you need the data;
> Why it's not creepy and is, in the grand scheme of things, fair.

It can be done in 1-5 sentences! **If your use-case is simple, don't overthink it.**

> Example: "We log failed login attempts with IP addresses and timestamps to detect brute-force attacks. // interest
> This is necessary to prevent unauthorized access to user accounts. We can't reasonably protect accounts without monitoring attack patterns. // necessity
> We don't use it to track users, and we only retain logs for 30 days — there is no way it can be used to harm them." // balancing

PS: If you find yourself explaining why users "shouldn't mind" your processing in more than 3 sentences, you probably need consent instead. Would your aunt be shocked to hear about it? Drop it.

... Aaaand: done!

> *That's it. You have a legal basis!!!*

Congratulations!

You are *officially* done with one of the hardest parts of the GDPR!

To recap this chapter: for your processing to be legal, **for each purpose you are using data, you need one Article 6 legal basis.** You get to pick between several, but in practice, it will most often be **consent, contract,** or **legitimate interest.**

So, what you need is a list. You take the one you **totally did make in the previous chapter (**one with your purposes), and next to each one, you jot a legal basis.

Profiling for ads? Consent. Debugging? Legitimate interest. Account creation? Contract.

There is no perfect system here. **Nobody really knows how to do it with a 100% accuracy.** I mean, you can read CJEU judgements and weep, and still not know.

But even trying gets you way, way closer to compliance than you can imagine.

Summary

Rule of thumb: you want to give people a real choice, and processing of data is totally optional? Consent. Make sure it's for that one, specific purpose, that there are no negative consequences if people say 'no', and that it's formulated in a simple way.

Your app would break without it? Contract. No need to do anything. (Apart from jotting that down somewhere, but that goes for everything.)

More useful to you than invasive to people, necessary to achieve something legal (like make money), reasonably expected by people? Legitimate interest. Write down your interest, which data you use, why it's necessary, and how it affects people - two sentences. Ideally, give them an option to opt-out. (For marketing stuff, you have to. For other stuff, say security or debugging, you do not – they can object, in which case you can consider to delete the data. But that is a whole another story.)

... Aaaand: I lied.

> There are *exceptions*, because of course.

So just to note: this is not *truly* the whole picture. Like, in *most* of the cases it is, but there are exceptions.

I can't cover them all here, but let's see..

First, **cookies** and tracking technologies. This is not purely regulated under the GDPR – there is another monstrosity called the 'ePrivacy Directive' of 2002. Under its Article 5, 'use of electronic communications networks to store information or to gain access to information stored in the terminal equipment of a subscriber or user' requires consent. The exception being 'any technical storage or access for the sole purpose of carrying out or facilitating the transmission of a communication over an electronic communications network. or as strictly necessary in order to provide an information society service explicitly requested by the subscriber or user'.

In human terms: if you are using cookies and other tracking technologies, **you need consent.** If the cookies are 'strictly necessary', i.e. your website/app would break without them, you do not. What is 'strictly necessary'? Nobody knows, but think important session cookies, authentication, payments.

FFS I *hate* the ePrivacy Directive and its idiotic 'terminal equipment', nay, me sayeth ARPANET, but hey, it's here.

Secondly, you normally need consent to send **marketing** e-mails and offers — or do any kind of direct electronic marketing, *unless* you are promoting your own goods or services which are similar to those the user has bought before. In other words, if user bought one game from you, and you have their e-mail, you can, in principle, send them promotional e-mails without consent, as long as you offer opt-out: you would rely on legitimate interest. Be careful, though: this tends to creep out people.

Last round of exceptions – special categories of data. Think health, religion, sexual orientation, stuff like that — you can't use contracts or legitimate interest, so it would have to be consent (or some kind of a really special scenario).

Now, bear in mind: this is what *the GDPR* says. App stores might have different, sometimes stricter policies. We will dive into those a bit later.

Your favorite time! Recap update time!

> **Chapter 1:** GDPR applies if you are established in the EU or trying to sell stuff/monitoring people in the EU. It applies whenever you decide people's data will be used in some way, or store/process it for others. 'Data' means anything, like literally anything, even a damn session id which exists for 3ms.

> **Chapter 2:** To know what you need to do under the GDPR, you must start with a list of specific reasons why you are using people's data. Each one is called a 'purpose'.
>
> To comply with the core rules: first, delete all the stuff you don't truly need, and make sure you set up a reasonable deletion routine per each purpose. Secondly, secure the data depending on how much of a risk it poses, don't be an ass, and give people proper info. Thirdly, be careful if re-using data collected for one purpose for something unrelated. Not necessarily in this order.

> **Chapter 3:** For each purpose you have, you need a legal basis – one of the six alternatives. One 'to let users log on', one to 'send marketing e-mails', and so on.
>
> Consent is an alternative most people talk about. But consent is tricky. *Nothing* negative can happen to a user if they say no (disabled feature), it can be withdrawn whenever, it must be granular, informed, active. Avoid when in doubt.
>
> Two important alternatives to know. First, if not using data breaks your app or service – use 'contractual necessity' instead. If what you are trying to do is reasonable, doesn't hurt users, they can expect such stuff to take place, go for 'legitimate interest' – document that this is sane and that you don't need consent. This does not always apply to cookies/marketing/sensitive data, but that's the gist.

TBC.

04

Privacy policies

Privacy policies

> *It does not have to suck.*

But it *does* have to wait.

So, you are ready to publish your app! Yaay! First things first, let's write a privacy policy, right!?

Wrong.

Privacy policy must come at the *end*. You cannot write a privacy policy until you've gone through your *purposes, principles* and *legal basis.*

Meaning, without having gone through your codebase, looked at the data you're using, and, crucially, reasons why, no lawyer, no AI, no deity can write one for you.

But if you are done with the steps so far? *Should be simple!* - Ish.

The starting point are GDPR Art. 13 and 14, which are so *utterly horrible*, I have to show them. I'll go with 14.

Behold.

Article 14

> *Or, what your privacy policy must contain. We'll translate. For now, just look at how -insane- it is!*

1.Where personal data relating to a data subject are collected from the data subject, the controller shall, at the time when personal data are obtained, provide the data subject with all of the following information:

a)the identity and the contact details of the controller and, where applicable, of the controller's representative;

b)the contact details of the data protection officer, where applicable;

c)the purposes of the processing for which the personal data are intended as well as the legal basis for the processing;

d)where the processing is based on point (f) of Article 6(1), the legitimate interests pursued by the controller or by a third party;

e)the recipients or categories of recipients of the personal data, if any;

f)where applicable, the fact that the controller intends to transfer personal data to a third country or international organisation and the existence or absence of an adequacy decision by the Commission, or in the case of transfers referred to in Article 46 or 47, or the second subparagraph of Article 49(1), reference to the appropriate or suitable safeguards and the means by which to obtain a copy of them or where they have been made available.

1. In addition to the information referred to in paragraph 1, the controller shall, at the time when personal data are obtained, provide the data subject with the following further information necessary to ensure fair and transparent processing:

a)the period for which the personal data will be stored, or if that is not possible, the criteria used to determine that period;

b)the existence of the right to request from the controller access to and rectification or erasure of personal data or restriction of processing concerning the data subject or to object to processing as well as the right to data portability;

c)where the processing is based on point (a) of Article 6(1) or point (a) of Article 9(2), the existence of the right to withdraw consent at any time, without affecting the lawfulness of processing based on consent before its withdrawal;

d)the right to lodge a complaint with a supervisory authority;

e)whether the provision of personal data is a statutory or contractual requirement, or a requirement necessary to enter into a contract, as well as whether the data subject is obliged to provide the personal data and of the possible consequences of failure to provide such data;

f)the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject.

1. Where the controller intends to further process the personal data for a purpose other than that for which the personal data were collected, the controller shall provide the data subject prior to that further processing with information on that other purpose and with any relevant further information as referred to in paragraph 2.
2. Paragraphs 1, 2 and 3 shall not apply where and insofar as the data subject already has the information.

Let's fix it!

Article 14, humanized

1. Where personal data relating to a data subject are collected from the data subject, the controller shall, at the time when personal data are obtained, provide the data subject with all of the following information:

a) the identity and the contact details of the controller and, where applicable, of the controller's representative;

// Who you are and how to get in touch.
*// "Hi! I'm Milos Novovic, and you can reach out at **milos.novovic@bi.no**"*

---

b) the contact details of the data protection officer, where applicable;

// If you have appointed a person who is a DPO, give a contact – if you are reading this guide, I assume no?
// "You can get in touch with our DPO at dpo@abc.ab."

---

c) the purposes of the processing for which the personal data are intended as well as the legal basis for the processing;

// Crucial! Your 'Milos-is-nagging-about-homework'-list of purposes from Chapter 2.
// "I use data to send marketing e-mails, to analyze crashes, to secure my app, to..."

---

d)where the processing is based on point (f) of Article 6(1), the legitimate interests pursued by the controller or by a third party;

// Decided to go with LI instead of consent for a purpose? Say what the interest is.
// "We process data to ensure system security."

---

e) the recipients or categories of recipients of the personal data, if any;

// Whom are you sharing the data with? Does not (usually) need to be granular to the level of individual companies, but it's a good idea to be as transparent as possible, especially given that you should already have this info available anyway.
// "We share your data with our cloud storage providers (Supabase), security incident logging databases, and, if you chose to use AI features, providers of those services (OpenAI, Google)."

---

f) where applicable, the fact that the controller intends to transfer personal data to a third country or international organisation and the existence or absence of an adequacy decision by the Commission, or in the case of transfers referred to in Article 46 or 47, or the second subparagraph of Article 49(1), reference to the appropriate or suitable safeguards and the means by which to obtain a copy of them or where they have been made available.

// Oh dear God, there is a whole drama behind this one, and I do not have the time to write a chapter on the nonsensical provisions on third-country transfers. But here's the reality: most services you're already using (Firebase, Supabase, OpenAI, etc.) have Standard Contractual Clauses in place — basically contracts saying "we'll handle EU data properly even outside the EU" — *that's* what you need to refer to.
// So.. unless you're dealing with high-risk data or sketchy providers, just check that your main services have these agreements (a quick Google search usually tells you), mention where you transfer data, and move on. This is way less important than getting your core purposes and legal basis right – do a deep dive into it if needed, but do the *basics first*. I honestly don't know how to phrase 'this is a requirement, but for most people not on the top of the list, but everyone talks about it, but it's not that risky in practice, but I am not advising you to break the law' differently.
// "We transfer your personal data to the United States and Canada on the basis of Standard Contractual Clauses. Email us for a copy."

... pt 2.

2. In addition to the information referred to in paragraph 1, the controller shall, at the time when personal data are obtained, provide the data subject with the following further information necessary to ensure fair and transparent processing:

a) the period for which the personal data will be stored, or if that is not possible, the criteria used to determine that period;

// When do you delete the data? We discussed deletion back in Chapter 2, the storage limitation principle – you delete it when you no longer need it. If you have a specific number of days, give it; if not, explain how you decide whether to delete.
// "We delete security logs 30 days after the last incident. Your purchase history is held as long as necessary to comply with applicable legislation. Your user account information is retained as long as you have an open account."

---

b) the existence of the right to request from the controller access to and rectification or erasure of personal data or restriction of processing concerning the data subject or to object to processing as well as the right to data portability;

// We haven't talked about people's rights in this guide, but basically, this just says you have to let them know they have this right. The analysis can be a bit tricky (sometimes they might *not*), but, to keep it simple, just state it.
// "You have the right to request access to your data, its correction or deletion, to object to its processing, or request it in a portable format."

---

c) where the processing is based on point (a) of Article 6(1) or point (a) of Article 9(2), the existence of the right to withdraw consent at any time, without affecting the lawfulness of processing based on consent before its withdrawal;

// Same.
// "Whenever you give us your consent, you are free to withdraw it at any time."

---

d) the right to lodge a complaint with a supervisory authority;

// Same.
// "You can always complain to a supervisory authority."

---

e) whether the provision of personal data is a statutory or contractual requirement, or a requirement necessary to enter into a contract, as well as whether the data subject is obliged to provide the personal data and of the possible consequences of failure to provide such data;

// This is the list we have made back in Chapter 3 – where we added Article 6 legal basis next to each purpose. Just integrate with info we've already covered.
// "We send you promotional offers based on your consent. We also analyze app crashes, on the basis of legitimate interest, to make sure our app runs smoothly. Since this is an image processing app, we process your images, so that we can fulfill our contract with you."

---

f) the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject.

// Do you have an app which makes legal decisions about people, or makes *really* important decisions about them – say, do they get a job or not? If yes, explain the criteria.
// "We do not use your data to make decisions about you which would create a legal or significant other effect."

Or a better fix:

1

Step 1:

Introduce yourself and your DPO (if you have one).

We are ABCD, a company registered in Norway. You can reach out to us, or our data protection officer at info@abcd.ab.

2

Step 2:

Get your list of purposes/legal bases and turn them into a sentence/two.

We use personal data to let people highlight legal texts, cross-sync them, pay for a 'pro' version, all of which we do so we can fulfill our contract and deliver you the service. We also send promotional emails if you consent to receiving them, and we monitor the app for irregular log-in attempts, in order to deliver a secure service, which is both in our and your interest.

3

Step 3:

Add specifics on storage periods, data sharing and international transfers.

We delete all your user data once your account is deleted. We delete your e-mail address 15 days after you unsubscribe from our promotional e-mails. Security logs are kept for 30 days, unless there is an active investigation pending.
We share your data with our cloud services provider, Amazon, and its sub-processors. Some of data processing takes place outside of the EU, on the basis of Standard Contractual Clauses, which you can contact us about.

4

Step 4:

Add boilerplate

You can always withdraw any consent you have given, and you can object to processing based on legitimate interest. You have the right to request access to your data, its correction or deletion. You can always complain about our data processing to a supervisory authority, but we'd appreciate a heads-up. We do not profile you as to make legal or significant decisions about you.

5

Step 5:

Profit.

Done!!

05

All good things...

Come to an end.

And so does this guide.

By now, you have an idea of whether the GDPR applies to you at all, which (sensible) stuff you have to do, when/how/whether to ask for consent, and how to write a privacy policy.

That is a huge achievement!!

And I hope that, if not outright *fun*, this was an OK ride.

Above everything else, I hope that you have seen that the GDPR is no monster — at all — and that the things it asks for are something you should — and even likely are — doing anyway. Despite its reputation, it is, at its heart, mostly about the common sense.

The reason why I, as a legal rather than an IT professional, love coding is because of that moment of pure joy when I grasp how something works — or find a cool new solution for an issue that's been bugging me for ages.[1] Like, WWDC is announced, and already in April, I am kinda pumped — have they finally made SwiftData usable? Can I use this in my app? Although I will, naturally, never release the said app.

Not to get all philosophical, but I think that kind of curiosity is in the human nature. A monkey who figured out he could use a stick to get a banana probably felt mighty good about himself. We're hardwired to be problem-solvers.

And that's why I reject the idea that law is boring, or complicated, or only for those who charge by the second. Thinking "hey, do I really need this data?", or — "is this creepy, or do people expect it" — is as much of an architectural and business question, as it is one of the GDPR.

This is not to say that the Regulation 2016/679 is an answer to all the problems in the world, but rather — that the same reasons that pushed you to look at the issues from those other perspectives, can be the same ones which will make the GDPR a perfectly fine journey. If nothing else, you *can* learn something new about your codebase and your business, hey!

To wrap things up: on the next page, I will give you a brief and *very* superficial overview of the things I did *not* cover in this guide. I think we are getting it to compile by now, but there are still some compiler warnings, so let me at least list the most important ones. After that, I will give you a rant on fines and the stuff which you should not buy. Then a few platform policies. Then I'm gonna sign off.

Fine, fine!

> HEY, YOU! STOP! You've decided to *breathe* today??

Thaaaat's it, a fine up to 20 million EUR, or 4% of your global annual turnover, whichever one is larger!

Well, *no*.

Let's get this straight: a fine is a penalty issued by a supervisory authority — one of the country where you are established; coordinated, in cases with international data flows, with others. Now, **a fine is only *one* tool at their disposal:** usually, as we have seen in practice, they **won't reach *straight* for it.**

Instead, what the GDPR lets them do, is notify you that they have spotted irregularities and ask you to correct them; publicly shame you (I kid you not); order you to stop transferring data; and so forth.

Most authorities *will*, in other words, choose to be reasonable about it — and not just out of being nice. One of the core legal principles we have is called *proportionality* — think how a police officer should not shoot me if I am jaywalking. I mean, he *can*, but no court would consider his action legal. The same goes here: unless you were doing something *truly* wrong, your first interaction is likely *not* going to be a fine.

And even if it is: **it is not likely to be 20 million EUR.**

How do they come up with an amount of a fine?

It seems rather like a 'make up a number and roll with it' kind of methodology in practice, but we do see *some* criteria under the GDPR.

> First — the nature of what you did, kind of data you've used, number and kind of users. An app for kids which profiles them and sells data to religious organizations so they can convert them? You can imagine. A to-do app developer with 200 users? Not so much. Secondly: was what you did intentional or accidental (the latter does not exclude liability but the former makes it worse); did you try to mitigate things which went wrong? Thirdly: did you tell the authorities that you infringed the law, or did they have to find out on their own *(right)*. Fourthly — what's your track record? Is this the sixth time you suffer the same kind of a data breach, caused by the same vulnerability? And so forth.

So, once again, **there is some clear room for reason in here.**

> **Being sued**
> People can sue if their GDPR rights are violated. They always have the right to sue in their home country.
> Don't confuse this one with fines — no, people *cannot* get up to 4% of global turnover; what they *can* get is *compensation* for whichever damage they have actually suffered. No penalties — the idea is to compensate for any harm actually suffered, [1] no matter how minimal.
> In other words, you had a breach? You can be sued, but they have to prove it hurt them (and yup, being afraid counts). But those amounts don't tend to be huge, and can't be 48 gazzilion dollars, US-style.

Stop throwing your money away! Please!

> *Do some legwork before you hire someone, assuming you do need to hire someone.*

Why are we all obsessed with fines?

Because they sell.

20 million is a number which looks impressive on slides. And scary. And important. Add the media headlines and the GDPR's horrible reputation, and you will have CFOs opening their wallets left and right.

Don't fall for it.

Don't let them sell you a useless project that will just bog you down for the next three decades as they ponder whether your processing of a pencil order is going to constitute a risk to right to someone's right to data portability in 2029. This is coming from an *academic*, mind you, so *yeah*, that's how bad it gets.

Also: come prepared.

You do *not* want someone to charge you for compiling a list of purposes for which you are using the data. Only *you* know what you need the data for. Or for deciding which features qualify as the core ones in your app (and thus need no consent, say). You'd be paying someone to literally *sit and wait* while you are, well, making yourself compliant.

Instead, seek advice once you have an overview of your goals and processes. Ideally, after you have done things we talked about here. Even if you get some wrong: at least start. Then have someone take a look and help you do a sanity check.

And also...

Be difficult.

Want to hire someone? Ask them whether you need consent for security logging or crash analytics. If they say 'yes', charge *them* for wasting your time. Write 'legitimate interest' on the bill. If you know better after three hours with this guide, they should, too. Unless they give you a really thoughtful answer. Or bring cookies.[1]

And for the love of God, *never* buy a 'this-will-make-you-GDPR-compliant'-tool.

No, it won't 'help you avoid the fines'. Sure, some of them may help you systematize your docs, but damn, why pay a premium when it can all be done with... WordPerfect? Or, for that matter, on a clay tablet? And done *better?*

FFS, if you still wanna buy a GDPR AI-SaaS product-compliance tool after reading this guide, hire *me*. (Likely still burning your money without reason, but at least I am more personable than FineExterminator 3000. Hopefully.)[2]

So: do the work, have some fun with it, and get some help when you hit a snag, not before!

Platform policies

So, this being a GDPR guide, we have covered... Well, *the GDPR.*

But much to its *personal* dismay and offense, the GDPR is only one part of the privacy picture.[1] There are countless other laws at play — depending on where you offer your app, whether you use some kind of an AI, whether you deal with kids, health stuff, or offer an 'online marketplace for sale and marketing of explosive precursors'.[2]

But, in addition to laws, there are also *platform* policies. To get your app on Apple's App Store, for example, you have to comply with *their* privacy guidelines — and not just the GDPR.

Which means: you could do everything by the book under the GDPR, and be in breach of Apple's policies. And conversely.

Now, I am *not,* repeat, *not,* in any possible way associated with Apple, I have never published an app, and I *do not know* how they enforce their policies in practice. Take what I say with *even more* salt than usually.

So.

When publishing your app, first, you need to have a privacy policy (see **guidelines 5.1**), and you need a Privacy Nutrition Label. This shouldn't be an issue — once you've completed the previous chapter, it should go just fine.

Secondly, there's the App Tracking Transparency (ATT): in brief, Apple requires you to ask users for *consent*, using the specific system dialog, *if* you are *tracking* them. Tracking is defined as "linking user or device data collected from your app with user or device data collected from other companies' apps, websites, or offline properties for targeted advertising or advertising measurement purposes", or "sharing user or device data with data brokers".

Two direct implications: you can't say 'legitimate interest holds up under the GDPR, I don't need consent' if your use-case involves cross-tracking or sharing data with brokers. Secondly, no matter how well-crafted your GDPR consent prompt might be, you need to use the ATT one. Which might lead to an odd situation where you have two different consent prompts. Ugh.

But.. the real devil is in the ecosystem: yup, all the third-party SDKs. Let's talk... Examples.

Firebase. It can come with ads. It can come without. Whether you need consent depends entirely on which *parts* you use and how you *configure* them.

> Say you use Crashlytics to help you fix bugs. It uses its own UUID to count affected users; it doesn't use IDFA and isn't linking data for ad purposes. Based on what I could find, at least. So, no tracking, no ATT prompt needed. Under the GDPR, bug-fixing is a textbook case for "legitimate interest".
>
> But say you import conversion events (like first_open or in_app_purchase) from Firebase into Google Ads to see which ad campaigns are driving installs. Or you export audiences you've created in Firebase (e.g., "users who completed level 5") to Google Ads to retarget them with ads in other apps.
>
> Would legitimate interest hold up under the GDPR? *Maybe,* but unlikely. But the question is moot, because ATT requires consent, anyway.

When deciding whether to fire ATT prompt, you have to be familiar with the SDK you're using and what it does with user data. Where does the data end up? Will it be used for some other purposes? Have I checked the default settings?

This is **not**, I stress, **not**, purely theoretical. In Denmark, a municipality got slapped with a fine b/c they used Google Ed suite on the laptops they gave to kids. Not an issue, I hear you say, and that was my first instinct, too: I rolled my eyes so hard I am still looking for my left eyeball. But. They actually had a point. The system made YouTube accounts for kids by default, with their full, real name. The kids didn't know that those accounts were public, started leaving comments, and... Yup. The whole mess. Why? Because someone didn't check the defaults.

Check the defaults.

To keep it simple: targeting ads/measuring ad performance/sharing data with ad brokers? ATT.

Before you say 'I am not': check what the SDKs you have imported do.

And then *very carefully* decide which ones to import.

(Say, you skip 'Firebase/Analytics', and use 'Firebase/AnalyticsWithoutAdIdSupport' instead.)

And you check the living Hell out of the default settings. And docs.

And you — you've guessed it — write it down. Pen and paper or .md.

The stuff I skipped

With the GDPR being long and complex, there are things I inevitably had to skip in this guide; things that you might be legitimately interested in. (See what I did there? I'm a riot. And you are amazing for having paid attention.)

The point being: here comes a list of things small and large which I have not had a chance to explore in meaningful depth. Give them a skim: in some future versions, I can turn them into mini-chapters, depending on what the people have to say.

Miscellaneous stuff

Data breaches

> Suffered a security breach? Breathe breathe breathe! It's fine.
> Here is what you need to do: first, act quickly to put out any fires. Secondly: you need to assess how serious the breach was. Minor? Just improve security measures in the future. Likely to result in 'risk to rights and freedoms'? You have to notify a supervisory authority within 72 hours. (They are usually a bit flexible on the deadline, but reach out. If you are outside of EU/EEA, you need a legal representative. Don't pay a fortune, though, they are glorified secretaries. Unless also your DPOs.) Even higher risk to rights/freedoms? You have to notify the affected people. An e-mail, in-app notification, common language. See Article 33/34.

Impact assesments

> Do you have a huge database, and are matching it over multiple datasets? Do you have tons of sensitive data, and/or offer services to kids, elderly, vulnerable people? You need to write a 'data protection impact assessment' — a doc which summarizes the risks to people, and lays down how you're gonna mitigate them. If the risks cannot be mitigated, you can reach out to the authorities, who should, in theory, help. See article 35.

Contracts and due diligence

> It's been way too long since we started, but you may recall that I said that most of this guide applies to data controllers — likely, you — i.e. those who decide how/why data will be processed. Those who decide nothing are called 'data processors'. Again, it has nothing to do with the tech stuff — who has the data — it's simply about who's making the calls. Now, if you are a controller, and you use a processor — say, your app uses Firebase — you have two obligations: first, you need to check if they follow the GDPR; secondly, you need to have a contract with them. Now, major services usually offer pretty good docs on their compliance, and will already have those contracts incorporated into ToS; but it's worth checking. In any case — less known vendor — run a search on them and data breaches/incidents. Re: major ones — watch out for one thing: do they use data for their own purposes? For instance, to improve an AI model? If so, you might be joint controllers, which can cause a world of hurt. You have to explain their processing to users as well, assess the consequences, the whole 9 yards.

People's rights

> Now, people's rights are important, and the GDPR is definitely empowering. But it is also, largely, reasonable. You *do* have to answer any requests from people ASAP, by latest in 30 days, or, if the request is super complex, within two additional months. You have to make it easy to submit requests, and should check the identity of people submitting them (within reason — sometimes a verified e-mail will do just fine).
> Things they can request?
> First, access to their data/a copy of it. Before you give it, you have to check that giving them access is not going to be a disaster to other people and interests (I ask for a copy of an internal investigation report while it's still pending; I ask for info which you cannot possibly give me without giving me access to your codebase). Secondly, they can ask to have their data deleted. Note: this right is actually very narrow, as it, basically, says that you have to delete the data when you're already processing it illegally (without a purpose, after people have withdrawn consent, without a legal basis, etc.). Thirdly, people can say that you don't have a legitimate interest where you claim one and ask you to stop using their data. If what they are objecting to is marketing (which you do without consent, right), you have to stop marketing to them — basically, the GDPR's opt-out for marketing. However, if your interest is something like security/usage analytics, you consider whether *this particular* person would be so harmed by it that you have to stop using *their* data; if not, you don't have to. Basically: a user can't ask you to turn off caching. Fourthly: if you make legal/very significant decisions about others in an automated way, you need a legal basis and they get human review. We also have data portability, accuracy, restriction of processing, but I'd file those under 'for later' for the most typical of apps.

Instead of a goodbye...

There!

Done!

The dragon has been... bored to death, a happy hero gets to go home!

Your GDPR work has come to an end!

... Not really. Not truly. It never ends, and that's kinda the point, as we've seen.

But by now, you are *practically almost compliant* and it did not take *that* long of a time!

... A plea for help

This is an ultra, super-duper-early alpha of the doc. I ramble. I need to update. I need to optimize.

And I *really really* cannot do it without you. So please: give me feedback & let me know what worked & did not work. You can reach out at **www.milos.no** or drop me a line at **milos.novovic@bi.no**. If this is to become a community resource which is *actually* useful & does some good in the world, it needs to be improved, and it can be improved only if I know what was unclear!

So: any (human) feedback welcome!

All the best,

Milos