URADICAL

POCKET REFERENCE

GDPR

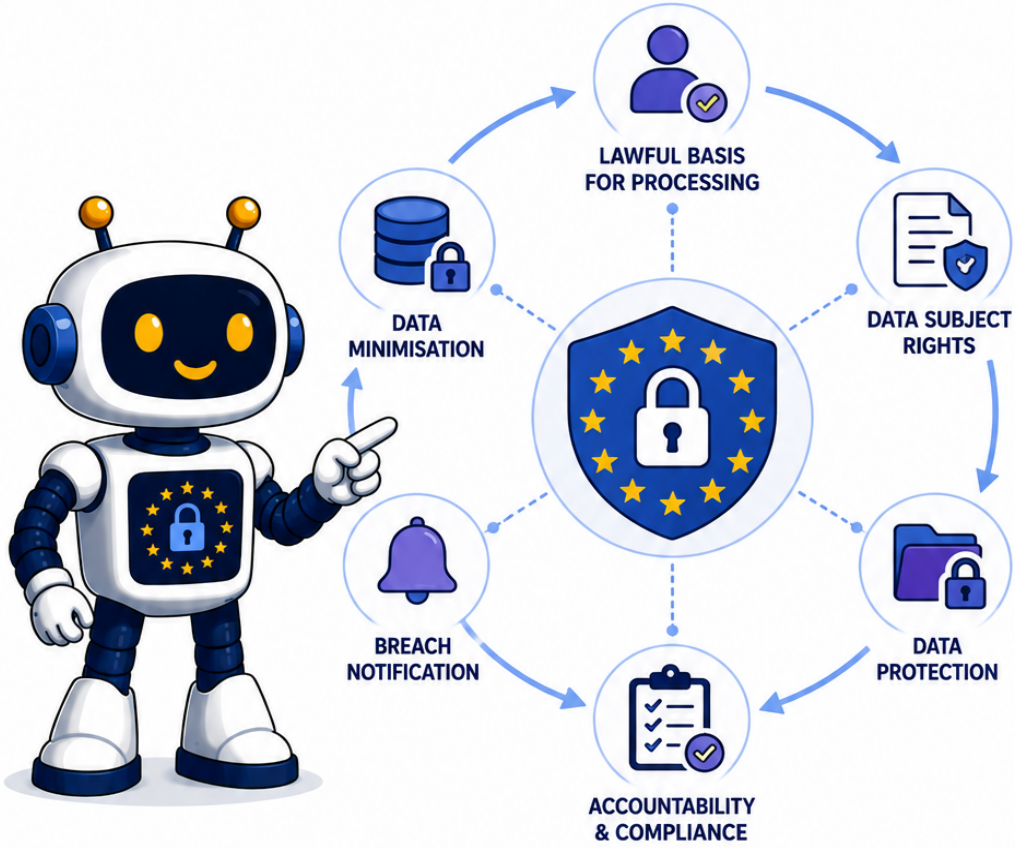

Alan Bradley

GDPR Pocket Reference for Developers

A concise, technology-agnostic reference for software developers and engineers

Alan Bradley

uradical.io

GDPR Pocket Reference for Developers

Alan Bradley

© 2026 uRadical

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

You are free to share and adapt this work for non-commercial purposes, as long as you give appropriate credit and distribute your contributions under the same license.

https://creativecommons.org/licenses/by-nc-sa/4.0/

Table of Contents

Introduction

This pocket reference is not legal advice. It is a practitioner's guide to understanding and applying GDPR principles in software development contexts. When in doubt, consult a qualified data protection professional.

Most developers meet GDPR at the worst possible moment: a legal team drops a compliance requirement into an active sprint, a client asks whether the system is "GDPR compliant" during a handover call, or a breach notification lands and nobody is sure what the 72-hour clock means. The regulation gets treated as someone else's problem until it isn't.

That framing is the problem. GDPR is not a legal wrapper applied around a finished system. It is a set of obligations that shape how systems must be designed — what data you collect, how long you keep it, what you must be able to do with it on demand, and what happens when something goes wrong. The decisions that determine compliance are made in design documents and schema reviews, not legal sign-offs.

This reference is written for developers who want to understand what the regulation actually requires, not just which checkboxes satisfy an auditor. The goal is to give you enough grounding to ask the right questions at the right stage — before the data model is locked, before the third-party SDK is integrated, before the consent flow is shipped.

It is deliberately technology-agnostic. The obligations apply whether you are building a Go microservice, a Rails monolith, or a serverless pipeline. The regulation does not care what language you use. It cares what you do with personal data and whether you can account for it.

Chapter 1: Scope and Key Definitions

What GDPR Covers

GDPR applies to the processing of **personal data** of individuals in the European Union. It applies regardless of where the organisation processing the data is based. If you are building a system in Belfast, San Francisco, or Singapore that handles data belonging to EU residents, GDPR applies to you.

The UK retains its own equivalent — **UK GDPR** — which is substantively identical following Brexit. For practical purposes, if your system is compliant with EU GDPR it will be compliant with UK GDPR, though you should verify this with a legal professional for your specific context.

Personal Data

Personal data is any information that relates to an identified or identifiable natural person. This is broader than most developers initially assume.

The following are personal data:

- Name, address, email address, phone number
- IP addresses and device identifiers
- Location data
- Cookie identifiers and browsing history
- Biometric data
- Pseudonymised data where re-identification is possible

The following are **not** personal data:

- Genuinely anonymised data where re-identification is not reasonably possible

- Data about companies or organisations (though data about sole traders may qualify)
- Data about deceased individuals (though member states may extend protections)

The key test is identifiability. If data can be combined with other data to identify a person — even indirectly — it is personal data.

Special Category Data

Certain categories of personal data attract heightened protection and stricter processing conditions:

- Racial or ethnic origin
- Political opinions
- Religious or philosophical beliefs
- Trade union membership
- Genetic data
- Biometric data used for identification purposes
- Health data
- Sexual orientation or sex life data

If your system touches any of these categories, the compliance burden is significantly higher. Processing special category data requires both a lawful basis under Article 6 and a separate condition under Article 9.

Data Subject

The individual whose personal data is being processed. In most systems, this is the end user.

Data Controller

The entity that determines the purposes and means of processing personal data. In most cases, this is your client or employer — the

organisation that decides *why* data is being collected and *what* is done with it. As a developer employed by or contracting for that organisation, you are building systems on behalf of the controller.

Data Processor

An entity that processes personal data on behalf of a controller. If you are a consultancy processing client data under instruction, you are likely a processor. Cloud infrastructure providers, analytics platforms, and email delivery services are typical processors.

The distinction matters because controllers and processors have different legal obligations and different liability exposure.

Processing

Processing covers virtually everything you do with data: collection, recording, storage, retrieval, use, disclosure, transmission, erasure. If your system touches personal data in any way, it is processing.

Chapter 2: Lawful Bases for Processing

Every act of processing personal data must have a lawful basis. There are six. You cannot process personal data without identifying which one applies before you begin. "We might need this later" is not a lawful basis.

1. Consent

The data subject has given clear, freely given, specific, informed, and unambiguous consent to the processing for a specific purpose.

Consent must be:

- **Freely given** — not bundled with terms of service, not a condition of using a service where opting out is not a genuine option
- **Specific** — given for a defined purpose, not a blanket agreement to any future processing
- **Informed** — the person must understand what they are consenting to
- **Unambiguous** — silence, pre-ticked boxes, and inactivity do not constitute consent
- **Withdrawable** — withdrawing consent must be as easy as giving it

Consent is often over-relied upon. If processing is genuinely necessary for a contract or legitimate interest, using consent creates compliance overhead with no benefit, and creates the obligation to stop processing when consent is withdrawn.

2. Contract

Processing is necessary for the performance of a contract with the data subject, or to take steps at their request prior to entering a contract.

This is the correct lawful basis for processing a customer's address to fulfil a delivery, or processing payment details to charge for a service. It does not extend to secondary uses of that data — analytics, marketing, or product improvement based on that same data require a separate basis.

3. Legal Obligation, Vital Interests, and Public Task

Three bases with narrower application:

Legal obligation — processing required to comply with law. Retaining employee payroll records for tax purposes, for example. The obligation must be specific and legally enforceable, not a general policy preference.

Vital interests — processing necessary to protect someone's life. Rarely applicable in software contexts outside emergency services.

Public task — processing necessary for a task carried out in the public interest or official authority. Primarily relevant to public sector systems. Not available to commercial organisations acting in a private capacity.

4. Legitimate Interests

Processing is necessary for the legitimate interests of the controller or a third party, unless those interests are overridden by the interests or rights of the data subject.

This is the most flexible basis and the most frequently misapplied. It requires a three-part test:

1. **Purpose test** — is there a genuine legitimate interest?
2. **Necessity test** — is processing necessary for that purpose?
3. **Balancing test** — does the interest override the data subject's rights and expectations?

Legitimate interests cannot be used for processing that data subjects would find unexpected or objectionable, or where the impact on them is significant. It cannot be used by public authorities acting in their public capacity.

Chapter 3: Data Subject Rights

Data subjects have eight rights under GDPR. Your systems must be capable of honouring them. This is not a legal formality — each right is a functional requirement. If the system cannot fulfil a right within the required timeframe, the controller is in breach. As the developer, you are the person who determines whether fulfilment is even possible.

Right of Access (Article 15)

A data subject may request a copy of all personal data held about them, along with information about how it is being processed, the purposes, the recipients, and the retention period. This is a Subject Access Request (SAR). You have **one month** to respond, extendable by two further months for complex requests.

Practically: can you query every data store in your system and produce a coherent export of everything held about a specific individual? That includes the primary database, logs, audit trails, analytics events, backups, and any third-party processors. If identifying all records for a given person requires manual intervention, you have a gap that will be painful to close under time pressure.

Right to Rectification (Article 16)

A data subject may request correction of inaccurate personal data, or completion of incomplete data. One month to respond.

Practically: is there a mechanism to update data held in all stores consistently? If a name change must be propagated across a primary database, a search index, an analytics platform, and an email service, can that be done reliably — or does it require manual updates across four systems with no audit trail?

Right to Erasure (Article 17)

Commonly called the right to be forgotten. A data subject may request deletion of their personal data where it is no longer necessary for its original purpose, where consent is withdrawn, where they have successfully objected, or where the data was unlawfully processed.

This right is not absolute. It does not apply where processing is necessary for legal compliance, the exercise of legal claims, or public interest tasks.

Practically: can you delete a specific individual's data across all stores — including logs, analytics, backups, and third-party processors — without breaking referential integrity or leaving orphaned records? Deletion is the hardest right to implement cleanly and the most frequently inadequate in practice. Systems designed without erasure in mind routinely store user identifiers in places that are never cleaned: application logs, error tracking services, analytics events, third-party CRMs. Audit this early.

Right to Restriction (Article 18)

A data subject may request that processing be restricted — data retained but not actively used — in certain circumstances: while accuracy is contested, while an objection is being assessed, or where processing is unlawful but the data subject prefers restriction to erasure.

Practically: does your system have a concept of a restricted account or record — one that is retained but excluded from all active processing, queries, and exports? This is a state that many systems do not model and cannot easily retrofit.

Right to Data Portability (Article 20)

Where processing is based on consent or contract and carried out by automated means, data subjects may request their data in a structured, commonly used, machine-readable format — typically JSON or CSV — and may request direct transmission to another controller.

Practically: can you generate a complete, structured export of a specific user's data on demand? This is distinct from a SAR response, which is about disclosure. Portability is about interoperability — the data must be in a format another system can consume, not a PDF summary.

Right to Object (Article 21)

Data subjects may object to processing based on legitimate interests or public task, and unconditionally to processing for direct marketing. An objection to direct marketing must be honoured immediately and without any balancing assessment. Other objections require the controller to demonstrate compelling legitimate grounds that override the data subject's interests.

Practically: is there a clear and accessible mechanism for users to opt out of marketing? Is it honoured in all channels — email, SMS, in-app — without requiring separate opt-outs per channel?

Rights Related to Automated Decision-Making (Article 22)

Data subjects have the right not to be subject to decisions based solely on automated processing — including profiling — that produce significant legal or similarly significant effects. Where such processing occurs, data subjects must be informed, must be able to request human review, and must be able to contest the decision.

Practically: if your system produces risk scores, eligibility decisions, content restrictions, or pricing that affects individual users and is generated entirely by an algorithm, Article 22 may apply. This right is frequently overlooked because teams do not classify their own systems as making "decisions" — but automated content moderation, loan eligibility engines, and personalisation systems that restrict access all fall within scope.

Right to be Informed (Articles 13 and 14)

Data subjects must be informed about the processing of their data at the point of collection (Article 13) or within one month if data is collected indirectly (Article 14). This is fulfilled through a privacy notice that accurately describes current processing activities.

Practically: is the privacy notice maintained as a living document, updated when processing activities change? A privacy notice written at launch that does not reflect the analytics, marketing, and third-party integrations added since is not compliant — regardless of whether it looks professional.

Chapter 4: Controller and Processor Obligations

Controllers

Controllers bear the primary compliance burden. They must:

- Identify and document lawful bases for all processing activities
- Maintain Records of Processing Activities (ROPA) — a written record of what data is processed, for what purpose, under what basis, with what retention period, and shared with which processors
- Implement appropriate technical and organisational measures
- Appoint a Data Protection Officer where required
- Conduct Data Protection Impact Assessments for high-risk processing
- Report personal data breaches to the supervisory authority within 72 hours
- Notify affected individuals of breaches without undue delay where the risk to them is high

Processors

Processors must:

- Process data only on documented instructions from the controller
- Ensure people processing the data are bound by confidentiality
- Implement appropriate security measures
- Not engage sub-processors without the controller's authorisation
- Assist the controller in honouring data subject rights
- Delete or return data at the end of the contract

Data Processing Agreements

Every controller-processor relationship must be governed by a written Data Processing Agreement (DPA). If you are a consultancy handling client data, you need a DPA in place with your client before processing begins. If you are using third-party services — cloud platforms, analytics tools, email providers — your client needs DPAs with those services, and you should verify they exist.

Data Protection Impact Assessments

A DPIA is required before beginning any processing that is likely to result in high risk to individuals. Indicators include:

- Systematic and extensive profiling
- Processing of special category data at scale
- Systematic monitoring of publicly accessible areas
- Use of new technologies
- Processing that prevents data subjects from exercising a right or using a service

A DPIA is not a one-time exercise. It should be reviewed when the nature of processing changes.

Chapter 5: Key Obligations

Records of Processing Activities

Controllers with more than 250 employees are explicitly required to maintain ROPA. In practice, all controllers should maintain them. A ROPA entry should capture:

- The name and contact details of the controller
- The purposes of processing
- The categories of data subjects and personal data
- Categories of recipients
- Third country transfers
- Retention periods
- A description of security measures

Data Retention

GDPR does not specify retention periods. It requires that personal data be kept no longer than necessary for the purpose for which it was collected. This means you must define retention periods, document them, and enforce them. Indefinite retention is not compliant.

Data Breach Notification

A personal data breach — any breach of security leading to accidental or unlawful destruction, loss, alteration, unauthorised disclosure, or access to personal data — must be reported to the supervisory authority within **72 hours** of becoming aware of it, unless the breach is unlikely to result in risk to individuals.

If the breach is likely to result in high risk to individuals, those individuals must also be notified directly and without undue delay.

In practice: does your organisation have a documented incident response procedure? Is it tested? Do developers know who to escalate to and when?

Privacy by Default and by Design

Article 25 requires that data protection is considered from the outset of system design, and that by default only the minimum necessary data is processed. This is not optional and it is not retrospective — it applies at the design stage.

Practically: privacy by design means the schema review is a compliance review. Every field added to a data model should have a stated purpose and a retention period. Privacy by default means that the most privacy-protective settings are active without user intervention — optional data collection is off by default, not on. A system that collects a user's date of birth because it might be useful for personalisation, with no stated purpose and no retention limit, fails both tests before a line of business logic is written.

International Transfers

Transferring personal data outside the UK or EEA is restricted unless the destination country has an adequacy decision, or appropriate safeguards are in place — typically Standard Contractual Clauses. This includes transfers to cloud infrastructure and SaaS tools hosted in the United States or other third countries.

Practically: the most common trap is assuming that a vendor's DPA is sufficient to cover the transfer. It is not — a DPA governs the processor relationship; it does not substitute for a lawful transfer mechanism. If you are integrating a US-based analytics platform, error tracking service, or CRM, you need to verify three things independently: that a DPA is in place, that a valid transfer mechanism exists (adequacy decision or SCCs), and that the vendor's sub-processors are also covered. Most

developers verify only the first and assume the rest is handled. It frequently is not.

Chapter 6: Current Enforcement Focus

GDPR enforcement has matured significantly since 2018. Supervisory authorities across Europe — and the ICO in the UK — have moved beyond basic compliance failures toward patterns of deliberate circumvention and emerging technology risks.

Dark Patterns in Consent Interfaces

Regulators are actively fining organisations for consent interfaces designed to manipulate rather than inform. Common patterns that attract scrutiny:

- **Pre-ticked boxes** for non-essential cookies or marketing preferences
- **Asymmetric design** — accepting all cookies is one click; rejecting requires multiple steps
- **Misleading language** — framing rejection as losing access to features
- **Consent walls** — making service access conditional on accepting non-essential processing
- **Endless layering** — burying controls in nested menus to discourage opting out

The principle is simple: consent must be as easy to withdraw as to give. If your consent interface makes rejection harder than acceptance, it is not compliant. The Irish Data Protection Commission, CNIL, and others have issued substantial fines specifically for dark pattern consent flows.

AI and LLM Processing

The use of large language models and AI processing pipelines introduces GDPR obligations that many development teams have not yet addressed:

Feeding personal data into third-party AI services — if personal data is submitted to an external LLM API, the provider of that API is a data processor. A DPA must be in place. The controller must have a lawful basis for the processing. Many organisations are submitting customer data to AI services without either.

AI-generated outputs about individuals — if a model generates information about a real, identifiable person, that output may constitute personal data, particularly if it can be attributed to them or used to make decisions about them.

Automated decision-making — Article 22 restrictions apply to AI systems making or significantly influencing decisions about individuals. Risk scoring, credit assessment, recruitment screening, and content moderation systems may all fall within scope.

Model training on personal data — using personal data to train or fine-tune a model requires a lawful basis. The model itself may retain personal data in a form that is difficult to audit or delete, creating right to erasure complications.

Regulators are increasingly focused on this area. Organisations cannot treat AI processing as exempt from the obligations that apply to any other processing activity.

Consent Manipulation and Tracking

Cookie consent enforcement has become aggressive. The CNIL has fined Google and Facebook hundreds of millions of euros for consent failures. The ICO has issued enforcement notices against major publishers. The pattern regulators look for: technically having a consent mechanism while designing it to fail.

Third-party tracking scripts loaded before consent is given, consent strings being set for purposes not consented to, and consent records that cannot demonstrate what a user actually agreed to are all active

enforcement targets.

If your system uses any third-party analytics, advertising, or tracking, the consent implementation deserves genuine scrutiny — not as a legal formality, but as a live enforcement risk.

Data Minimisation Failures at Scale

Enforcement is also targeting organisations that collect excessive data relative to their stated purpose — particularly where that data is then retained indefinitely or transferred to third parties. The combination of broad collection, long retention, and opaque third-party sharing has attracted significant fines.

The question regulators now ask is not just "do you have a privacy notice?" but "can you demonstrate that the data you hold is proportionate to the purpose you stated?"

Appendix A: New Project Checklist

Use this before writing a line of code. These are the decisions that cannot be retrofitted cheaply.

Scope and Lawful Basis

- [ ] Have you identified all categories of personal data the system will collect or process?
- [ ] Have you identified whether any special category data is involved?
- [ ] Have you identified the lawful basis for each processing activity?
- [ ] Is that lawful basis documented?
- [ ] If relying on legitimate interests, has a Legitimate Interests Assessment been completed?
- [ ] If relying on consent, is the consent mechanism freely given, specific, informed, and unambiguous?

Data Minimisation

- [ ] Is each data field collected genuinely necessary for the stated purpose?
- [ ] Have fields that are "nice to have" or "might be useful later" been removed from scope?
- [ ] Is the retention period for each category of data defined and documented?

Architecture and Third Parties

- [ ] Have all third-party services that will receive personal data been identified?
- [ ] Is a Data Processing Agreement in place with each processor?

- [ ] If any processor is outside the UK or EEA, is there an adequacy decision or are SCCs in place?
- [ ] Have you verified that third-party SDKs or analytics tools do not collect data outside consent scope?

Rights and Obligations

- [ ] Is there a mechanism to fulfil Subject Access Requests?
- [ ] Is there a mechanism to fulfil erasure requests across all data stores?
- [ ] Is there a mechanism to fulfil rectification requests?
- [ ] Is a privacy notice in place covering all processing activities?
- [ ] Is there a documented data breach response procedure?

High-Risk Processing

- [ ] Does the system involve profiling, automated decision-making, or special category data at scale?
- [ ] If yes, has a Data Protection Impact Assessment been completed?
- [ ] Has the DPO (if applicable) been consulted?

Appendix B: Audit and Review Checklist

Use this for existing systems, inherited codebases, and client handovers.

Data Inventory

- [ ] Can you produce a complete list of all personal data held in the system?
- [ ] Is a Records of Processing Activities document current and accurate?
- [ ] Are data retention periods defined and actively enforced?
- [ ] Is data being held beyond its defined retention period?

Lawful Basis and Consent

- [ ] Is a lawful basis documented for each processing activity?
- [ ] If consent is relied upon, are consent records stored and auditable?
- [ ] Does the consent record capture what the user consented to, when, and through what interface?
- [ ] Is the withdrawal mechanism functional and as easy to use as the consent mechanism?

Data Subject Rights

- [ ] Is a Subject Access Request process documented and tested?
- [ ] Has an erasure request ever been fulfilled end-to-end, including logs, backups, and third-party processors?
- [ ] Is there a process for handling rectification and restriction requests?
- [ ] Is there a process for responding to portability requests?

Third Parties and Transfers

- [ ] Are Data Processing Agreements in place with all processors?
- [ ] Are any processors outside the UK or EEA without an adequacy decision or SCCs?
- [ ] Have third-party tracking or analytics scripts been audited against the consent scope?

Incident Response

- [ ] Is there a documented data breach response procedure?
- [ ] Does the procedure include the 72-hour supervisory authority notification requirement?
- [ ] Has the procedure been tested or rehearsed?
- [ ] Is it clear who is responsible for making the notification decision?

Consent Interface

- [ ] Does the consent interface give equal prominence to accept and reject options?
- [ ] Are non-essential cookies and tracking disabled by default until consent is given?
- [ ] Is consent granular — can a user accept some purposes and not others?
- [ ] Does the privacy notice accurately reflect current processing activities?

Further Reading

- **GDPR full text** — eur-lex.europa.eu/eli/reg/2016/679
- **ICO guidance** — ico.org.uk/for-organisations
- **CNIL GDPR Developer Guide** — github.com/LINCnil/GDPR-Developer-Guide
- **EDPB guidelines** — edpb.europa.eu/our-work-tools/general-guidance