GDPR Compliance

A Practitioner's Field Guide — Volume 1: The Law and the Compliance Machine

Sylvan Press

Sylvan Press

Field-tested, plain-English references for the people who do the work and stand behind their call.

GDPR Compliance

A Practitioner's Field Guide, Volume 1: The Law and the Compliance Machine

Published by Sylvan Press, the security imprint of Sylvan Assurance, LLC, Vermont, USA.

ISBN (paperback): 9798951893154
ISBN (hardcover): 9798951893161
ISBN (ebook): 9798951893178

First edition, 2026.

Disclaimer. This book provides general guidance and recommended security and compliance practices, not legal, financial, or professional advice, and it does not create any advisory or consulting relationship between the author, the publisher, and the reader. Every recommendation is optional: a widely accepted way to reduce a common risk, and the reader decides which to adopt, adapt, or decline. Following this guidance reduces exposure to common risks but does not guarantee any particular outcome and does not prevent any particular breach, incident, or loss. Where this book refers to laws, regulations, standards, or commercial products, it does so for general awareness only; confirm specific obligations with qualified counsel, and note that product references are illustrative and do not constitute endorsement.

Sylvan Press, sylvanassurance.com

Contents

Preface

It is a Tuesday in March, and a freelance editor in Edinburgh is reading a vendor questionnaire from a German academic press. Twenty-two pages. It asks her to name the lawful basis on which she processes interview transcripts, to describe her sub-processor arrangements, to declare a retention period for each category of personal data she holds, and to confirm that her international transfers rely on a valid mechanism under the General Data Protection Regulation. She has been in business four years. She is the entire business. Until this morning she had taken it as read that GDPR is something multinationals worry about.

She is wrong; it applies. The German press will not be the last client to ask, and the questionnaire is going to take her a fortnight to answer properly. She is also not alone. Small operators, catching-up vendors, EU mid-market businesses, and multinationals all go through a version of that same fortnight. The size of the organisation changes; the shape of the moment does not. Someone realises that a regime they had filed under "not my problem" is in fact their problem, and they begin the slow work of building a programme around it.

This is a book for the people who have to do that work. There is no shortage of writing on GDPR-the-regulation, legal commentaries, article-by-article annotations, regulator guidance. Most of it is excellent on what the Regulation says, and very little of it is written for the practitioner who has to run a privacy programme on a Wednesday

morning: who needs to know what a Records of Processing Activities entry actually contains, when a Data Protection Impact Assessment is mandatory rather than merely advisable, how to answer a supervisory authority that has opened a routine enquiry, and what a defensible breach assessment looks like inside the seventy-two-hour clock of Article 33. The gap between knowing what GDPR says and operating a GDPR programme is wider than most operators expect, and the cost of crossing it the slow way is high, measured in regulator letters, customer churn, legal fees, and the quiet exhaustion of operators improvising the same answers for years because nobody wrote them down.

The frameworks here come from more than two decades of professional security and privacy practice inside a top-25 global company: including product security incident response, and network security and incident response before that. Privacy and product security overlap more than the boundary between the professions suggests: knowing which clock matters, what an evidentiary record looks like, and how to talk to a regulator without making things worse all transfer cleanly. This is a practitioner field guide, not a treatise. Where it cites an article, it cites it so you can find it; where it describes a practice, it describes it so you can do it. The aim is a working programme, not a perfect one: one you can defend on a Wednesday morning when a supervisory authority calls, and one your operators can maintain on a Thursday afternoon when the calendar is already full.

Four recurring organisations carry the worked examples, because a defensible answer at one scale is often the wrong answer at another: Linden Editorial, a one-person consultancy; Northwind Analytics, a US software vendor mid-rebuild; Heliotrope Innovaties, a Utrecht systems integrator with a mature privacy practice; and Argent Technologies, a 21,000-person multinational under the one-stop-shop mechanism. Wherever a chapter introduces a discipline, it shows that discipline at each of the four scales. One of them is close enough to your own to be worth reading first.

How to Use This Book

A reference book is only as useful as the path the reader takes through it, and the right path through this one depends very much on why you have opened it. The chapters are sequenced for a reader new to operational GDPR work, but very few practitioners actually read the book in chapter order. What follows is a short orientation, organised by the situation that brought you here.

You are new to GDPR and need to build a basic working understanding. Read Part I in order. Chapters 1 through 6 cover the foundations: what the Regulation is and what it is not, the controller-and-processor distinction that determines almost every downstream obligation, the six lawful bases for processing, the categories of personal data and the special-category data that carries heightened obligations, the data-subject rights regime, and the Records of Processing Activities (RoPA) discipline that is the operational backbone of any defensible programme. By the end of Part I you will have enough mental model to read the rest of the book in any order you choose.

You have inherited a privacy programme and need to work out where it stands. Start with the *GDPR Maturity Rubric*: Appendix B of Volume 3, the Operator's Workbook. Score your programme honestly against the four tiers (small-operator, catching-up vendor, operationally mature, multinational) and read the corresponding tier chapter in *Breach, Enforcement, and Operations* first. The rubric will surface the two or three disciplines where the gap between current and target

is widest, and those disciplines have dedicated chapters in Part II of this volume. Most inherited programmes have a recognisable failure shape, and naming it early saves months of generalised remediation effort.

You are a small operator who has just received a vendor questionnaire and is wondering whether GDPR applies to you. Read Chapter 1 for the framing. Read Chapter 2 for the controller-or-processor question that will determine most of your answers. Read the small-operator tier chapter of *Breach, Enforcement, and Operations* for what minimum-viable compliance actually looks like at the single-practitioner scale. Then return to whichever discipline chapters the questionnaire forces you into. You almost certainly do not need a Data Protection Officer (DPO). You almost certainly do need a RoPA. Chapter 8 explains the first; Chapter 6 explains the second.

You are a non-EU vendor selling into the EU and your privacy programme has been a procurement-questionnaire-response programme. Read the catching-up-vendor tier chapter of *Breach, Enforcement, and Operations* first; it is written for you. Then read Chapters 9 and 10, international transfers and vendor management, because those are the two disciplines where catching-up vendors most often discover that their existing arrangements do not stand up to scrutiny. Then read Chapter 6 (RoPA) here, and the opening chapter of *Breach, Enforcement, and Operations* on the Article 33 breach obligations; those two will give you the operational backbone you have been missing.

You are an in-house DPO or privacy professional in an EU-headquartered organisation. You will likely skim Part I and read Part II properly. Chapter 8 on the DPO function may be useful as a sanity check; the supervisory-authority engagement chapter of *Breach, Enforcement, and Operations* is the one readers in your position most often find hardest to see written down honestly elsewhere; that volume's chapter on the annual compliance cycle gives you a calendar you can adapt.

You are part of a multinational privacy office. Read the multinational tier chapter of *Breach, Enforcement, and Operations*, then Chapter 9 of this book together with its divergence and enforcement chapters: international transfers, UK GDPR and European Economic Area (EEA)

divergence, and the fines-and-enforcement landscape under the European Data Protection Board (EDPB). That guide's supervisory-authority engagement chapter is useful even at scale, because the patterns of constructive regulator interaction do not change as the organisation grows. The Workbook tier-3 templates assume your scale of operation.

You are coordinating a live breach and the clock in Article 33 is running. Close this book. Open the Workbook to the seventy-two-hour breach-assessment worksheet. Come back to the breach-window chapter of *Breach, Enforcement, and Operations* afterwards, when the notification decision is closed and you want to understand why a particular step mattered. The Workbook is built for the live moment; the Field Guide is built for the reflection that comes after.

You work for a supervisory authority or in policy. The chapters of *Breach, Enforcement, and Operations* on supervisory engagement, enforcement, and the multinational tier describe the operational reality that supervisory engagement and enforcement land on. The book is written for practitioners but should be readable as a window into how the rules show up in the work: particularly in the gap between what a small operator can plausibly maintain and what a multinational can be reasonably expected to demonstrate.

A note on the four-tier framing, because it shapes how almost every chapter is written. Each substantive chapter introduces a discipline at the conceptual level. The same chapter then shows what the discipline looks like at each of the four recurring organisational tiers: Linden Editorial, Northwind Analytics, Heliotrope Innovaties, Argent Technologies. The point of the tiering is to give the reader a way to locate themselves. If your organisation looks more like Northwind than like Argent, the Northwind treatment of a discipline is what you should be aiming for first. The Argent treatment is where you would expect to land in three to five years if you continue to scale. The Workbook's templates are tiered the same way, with explicit small-operator, catching-up-vendor, operationally-mature, and multinational variants for each artefact. The maturity rubric, Appendix B of the Workbook volume, is the cross-reference that ties the tiers to specific observable practices.

Two cautions about the tiering. First, the four tiers are not a ladder

you are obliged to climb in order. Plenty of small operators run defensible programmes at their scale and never need to grow into the Heliotrope shape; the Linden treatment is the right destination for them, not a staging post on the way somewhere else. Second, the tiering is not a maturity hierarchy in the sense that the multinational treatment is "better" than the small-operator treatment. Each tier has practices that work at its scale and would be wasteful or unworkable at another. Pretending you can run an Argent-style global privacy office at Linden's headcount is one of the most reliable ways to burn out a sole practitioner; pretending an Argent organisation can get away with a Linden-style RoPA is one of the most reliable ways to attract supervisory attention.

A final note. The Workbook's reference appendices are not optional reading. The glossary (Volume 3, Appendix C) keeps terminology consistent across chapters that touch overlapping topics: *controller*, *processor*, *joint controller*, *supervisory authority*, *lead authority*, *transfer mechanism* and so on are used in the Regulation's own senses, and the glossary records those senses for ease of reference. The maturity rubric (Volume 3, Appendix B) is referenced repeatedly in the body of this book and is one of the Workbook's most-used artefacts. If a chapter assumes you have skimmed one of them, the chapter will say so.

If you are not in any of the situations above, read Chapter 1 next. It is short and it sets the frame for everything that follows.

About This Book

GDPR Compliance is a three-volume set, and this volume is the first of the three. This book is organised in three parts, designed to be navigable both linearly and by topic, which is why the reading paths in the previous section move around so freely.

Part I: Foundations (Chapters 1 through 6). What GDPR is and is not; the controller-and-processor distinction that determines almost every downstream obligation; the six lawful bases under Article 6; the categories of personal data and the heightened obligations attached to special-category data under Article 9; the data-subject rights regime; and the Records of Processing Activities discipline. Part I is the part of the set that a reader new to operational GDPR work should read first; everything that follows (in this volume and the next) assumes the framing established here.

Part II: Core Disciplines (Chapters 7 through 10). Data Protection Impact Assessments, the DPO function, international transfers under Chapter V, and vendor and sub-processor management. Each chapter in Part II treats a discipline at depth and shows what it looks like at each of the four organisational tiers introduced in the Preface.

Part III: The Law Applied (Chapters 11 through 15). The data map in practice, consent surfaces and dark patterns, cookies and the ePrivacy boundary, privacy by design for builders, and marketing, profiling, and automated decisions. Part III takes the machine built in Parts I and II to the places where personal data actually lives.

The rest of the set: **Breach, Enforcement, and Operations** carries the regulatory landscape (the seventy-two-hour breach window, supervisory-authority engagement, UK and EEA divergence, fines and the EDPB), the three compliance-by-tier portraits (Linden Editorial for small operators, Northwind and Heliotrope between them for mid-market, Argent for multinationals), GDPR and AI systems, the annual compliance cycle, and the closing maturity chapter, with the supervisory-authority engagement templates riding along as its Appendix A. **Volume 3: the Operator's Workbook** holds the working artefacts: the DPIA set, the RoPA structure, agreement clauses, breach-notification templates, subject-rights responses, vendor due diligence, the template finder, the maturity rubric (Appendix B), and the glossary (Appendix C). Each volume stands alone. Each is modestly priced in ebook.

A note on what the book does not cover. It is not a survey of national implementing legislation across the EU Member States. For that you need either regulator-by-regulator guidance from each supervisory authority or a comparative-law resource. It is not an annotated commentary on the Regulation itself. The published commentaries are excellent and there is no point in duplicating them. The book does not cover ePrivacy, the EU Charter of Fundamental Rights, or the GDPR's interaction with the Law Enforcement Directive. It does not cover the sector-specific data regimes for health-research, telecommunications, and financial-services that sit alongside GDPR with their own rules. It touches the EU AI Act only at the intersections with GDPR, not as a regime in its own right. Where the book stops, it stops because the topic has been covered well elsewhere or because it sits outside the operational privacy practitioner's day job. The Workbook footnotes a short reading list of resources for the topics the book deliberately leaves alone.

Defensibility statement

The book frames its recommendations as practices that reduce regulatory and operational exposure in real GDPR programmes, not as actions that guarantee compliance with the Regulation or with any specific obligation under it. Compliance with GDPR is determined by the

Regulation, by the relevant national implementing legislation, by binding decisions of supervisory authorities and the European Data Protection Board, and by the courts; it is not determined by adherence to any one book's recommendations, including this one. Where the book describes a practice as defensible, the description means that the practice survives scrutiny in real operations: not that the practice is guaranteed to do so in your circumstances.

This is not legal advice

Nothing in this book is legal advice, and nothing in it should be treated as a substitute for legal advice. The author is a security and privacy practitioner, not a qualified lawyer in any jurisdiction. The book is written to help operators run defensible programmes; the specific obligations that apply to your organisation, in your jurisdiction, under your contracts, with respect to your processing activities, are for qualified counsel to confirm. Where the book describes a deadline, a notification obligation, a fine, a transfer mechanism, an article of the Regulation, or a regulator's enforcement posture, the description is intended for general orientation and may not reflect the current state of the law at the moment you read it. Confirm specifics with counsel before acting on them.

Framework as of date

The substantive content of this book reflects the state of the GDPR framework as of mid-2026. That includes the Regulation itself, adopted in 2016 and applied from May 2018. It includes the UK GDPR as retained and amended in UK law. It includes the most recent guidance from the European Data Protection Board. It includes the post-*Schrems II* transfer regime, with the EU-US Data Privacy Framework as it stood at time of writing. It includes the early operational guidance on the intersection between GDPR and the EU AI Act.

The regulatory landscape moves. Supervisory-authority guidance is updated continually. Court decisions reshape the operational reality of specific obligations. The EU AI Act and other adjacent regimes interact with GDPR in ways still being worked out by regulators and

by practitioners. The Digital Services Act, the Digital Markets Act, the Data Act, the NIS2 Directive, and the Cyber Resilience Act are the most consequential of those adjacent regimes. Treat the dates and the specifics in this book as a starting point for your own current research, not as a settled and timeless reference.

The three volumes

GDPR Compliance ships as three volumes. The two narrative volumes describe the practices; the Workbook volume holds the artefacts that go with them. Where the Field Guide explains the judgement behind a Records of Processing Activities entry, the Workbook supplies the template for the entry itself: in four tier-matched variants. Where the Field Guide explains how a defensible seventy-two-hour breach assessment proceeds, the Workbook supplies the assessment worksheet, the regulator-notification skeleton, and the data-subject communication template. The two volumes share the same recurring organisations, the same maturity rubric, and the same defensibility framing. You can buy and use any volume alone; readers who keep the set will find that each volume sends them to the others often enough that keeping all three close is the practical default.

Conventions used in this book

A few conventions are worth flagging. Acronyms are spelled out at first use in each major section of the book; readers skim and jump between chapters, and a single first-use spell-out at page one is not enough. The Request for Comments 2119 (RFC 2119) keywords MUST, SHOULD, and MAY appear in capitalised form only where the book is quoting or modelling policy excerpts. Elsewhere the book uses ordinary prose. Articles of GDPR are cited by number and short title at first use within a chapter: for example, *Article 6* (*Lawfulness of processing*) or *Article 33* (*Notification of a personal data breach to the supervisory authority*). Subsequent references within the same chapter use the bare article number. Single quotes appear inside double quotes. UK English is used throughout: *organisation*, *recognised*, *behaviour*, *programme*, *centre*, *catalogue*. Where the book recommends a practice, the recommendation is framed in the defensibility language described above.

Chapter 1 begins on the next page.

1 What GDPR Actually Is

Mira Linden is at her kitchen table in Edinburgh, on her second coffee, staring at a customer questionnaire from a German academic press.

The questionnaire is twenty-two pages long. It arrived as a PDF attached to an otherwise polite email from a procurement officer in Munich whose English is better than the English of most of Mira's clients. The press wants to engage Linden Editorial, Mira's one-person consultancy in Edinburgh, to copy edit a forthcoming monograph series on environmental economics. Before the engagement can begin, Mira is required to complete the questionnaire, sign an enclosed data processing agreement, and return both documents within ten working days.

The questionnaire opens with a section labelled *Articles 28 and 32: Processor Obligations*. The first question asks whether Linden Editorial maintains Records of Processing Activities under Article 30. The second asks for the name and contact details of Linden Editorial's Data Protection Officer. The third asks for a description of the technical and organisational measures in place to protect personal data. The fourth asks whether Linden Editorial transfers personal data to any third country outside the European Economic Area and, if so, on what legal basis.

Mira reads the questionnaire twice. She is not a lawyer. She has been running Linden Editorial for four years, after eleven years as a managing editor at an academic press in London. She knows manuscripts. She knows authors. She knows the rhythms of a peer-review cycle and

the small humiliations of arguing with a typesetter about a hanging indent. She does not know what a Record of Processing Activities is supposed to look like, particularly for a business that consists of one person and a MacBook.

She reads the questionnaire a third time. Then she opens a new tab and types *does GDPR apply to me* into the search bar, which is approximately where most people's compliance journey begins.

A lot of people start there. The story this book and its companion volumes tell, the story of the whole thing, is what happens between that search and a programme that actually works.

The shape of the regulation

The General Data Protection Regulation, Regulation 2016/679 of the European Parliament and of the Council, to give it its full name, was adopted on 27 April 2016, entered into force on 24 May 2016, and applied from 25 May 2018. It replaced the 1995 Data Protection Directive, which had governed European data protection law for the previous two decades and which had become, by general agreement, badly out of date. The Directive had been written in an era when most personal data lived on paper, in filing cabinets, in offices that the data subject could in principle drive to. The Regulation was written for a different era. An era in which personal data lives in a cloud provider's data centre on the other side of the world. An era in which it is processed by systems the data subject has never heard of, on behalf of organisations whose names they would not recognise.

Two structural points about the Regulation are worth establishing at the outset, because everything else in this book builds on them.

The first is that it is a **Regulation** rather than a **Directive**. The distinction matters. A Directive sets out objectives that EU Member States are required to achieve through their own national legislation, with each Member State retaining considerable latitude over how to do so. A Regulation, by contrast, applies directly in every Member State as it stands, without national transposition. When Article 33 of the Regulation says that a personal data breach must be notified to the competent supervisory authority within 72 hours, that is the rule in Stockholm,

in Lisbon, in Tallinn, in Dublin, and in every other corner of the Union. There are some Member State derogations (the digital age of consent under Article 8, for example, varies between thirteen and sixteen depending on the country) but the core regime is harmonised in a way that the Directive never was.

The second is that the Regulation has an **extraterritorial reach** that catches more organisations than people expect. Article 3 says that the Regulation applies to processing carried out in the context of an establishment in the Union, regardless of where the processing itself takes place. It also applies to processing carried out by a controller or processor *not* established in the Union, where the processing relates to offering goods or services to data subjects in the Union, or to monitoring their behaviour within the Union. A Texas-based Software-as-a-Service (SaaS) vendor selling to a customer in Hamburg is within scope. A consultancy in Edinburgh that processes manuscripts containing the personal data of research subjects in Berlin is within scope. The geographic centre of gravity of the organisation is not the test. What the organisation does, and to whom, is the test.

The United Kingdom left the European Union on 31 January 2020. Since then, the EU GDPR has continued to apply to UK organisations whenever they fall within the Article 3 extraterritorial scope, and a domestic version known as the UK GDPR, sitting alongside the Data Protection Act 2018, applies to UK-resident processing. The two regimes share a structure and most of their operative text, but they are no longer interchangeable. On 5 February 2026 the data-protection provisions of the Data (Use and Access) Act 2025 came into force, and they changed several things a practitioner relies on day to day: the clock on a subject access request, the consent rules for some analytics cookies, the treatment of solely automated decisions, and the test applied to international transfers. Where this book gives an answer that differs for a UK reader, it says so in a short note headed *In the UK*. Treat the two regimes as substantively aligned in structure and materially divergent in operation. The UK Information Commissioner's Office (ICO) enforces the UK version; the relevant EU Member State supervisory authority enforces the EU version. Mira, in Edinburgh, has to think about both: the UK GDPR for her UK-resident clients and her own op-

erations, and the EU GDPR for the German press's manuscripts and any other EU personal data that passes through her hands. This dual exposure is not unusual. It is the default for most UK-based operators with any kind of European client base.

It is worth pausing briefly on the architecture of the Regulation itself, because the structure quietly governs how the chapters of this book are organised. The Regulation runs to 99 articles, grouped into eleven chapters, prefaced by 173 recitals. The recitals are not operative law in the same way that the articles are; they are the explanatory text that sits in front of the operative provisions and gives them their interpretive frame. Supervisory authorities and the courts treat the recitals as authoritative guidance on how the articles are intended to be read. A practitioner who reads only the articles, and not the recitals that introduce them, will routinely miss the reasoning that decides hard cases.

The articles themselves are grouped into recognisable territory. Chapter I is the scope and definitions. Chapter II is the data-processing principles and the lawful bases. Chapter III is the rights of the data subject. Chapter IV is the obligations of controllers and processors. Chapter V governs international transfers. Chapters VI through VIII deal with supervisory authorities, the European Data Protection Board (EDPB), cooperation between authorities, remedies, liability, and penalties. The remaining chapters cover specific situations, delegated acts, final provisions. The arc of the book you are reading follows roughly the same arc. We will spend more time on Chapter IV (the operational obligations) than the Regulation's relative pagination would suggest. That is because Chapter IV is where most of an operator's working life is actually spent.

This is general information about the structure of the regime. It is not legal advice. Specific situations require specific advice from a qualified data protection lawyer. That disclaimer is repeated at intervals throughout this book, partly because it is true and partly because the regulatory subject matter rewards a degree of caution that goes beyond what the security or product-development books in this same series require.

Who this book is for

Most GDPR books are written for lawyers, by lawyers, and they do a perfectly good job of explaining the Regulation as a body of law. There is no shortage of those books. If you want a clause-by-clause commentary on Article 28 processor obligations, with a survey of the EDPB guidance and the most relevant Court of Justice decisions, the literature is rich and you should read it.

This book is not that book.

This book is for the person who has to *run* the programme. The in-house counsel at a mid-market vendor who has been handed the privacy portfolio because nobody else wanted it. The fractional Data Protection Officer (DPO) covering three small businesses on a Tuesday afternoon. The compliance lead at a SaaS company whose Series B due diligence has just surfaced the gap between the procurement-questionnaire answers and the operational reality. The security architect whose Chief Information Security Officer (CISO) has just said *figure out what we need to do for GDPR* and who is now looking at the regulation text with the dawning recognition that legal interpretation is not, in fact, what they trained for. The sole practitioner like Mira, working through the implications of one German questionnaire and one suddenly-load-bearing definition of *processor*.

The people in those roles do not need another commentary on the Regulation. They need a working understanding of how the Regulation actually lands when it meets a real organisation; and what to do about the parts that land hard.

The reader this book has in mind has the following in common with Mira, or with the people introduced shortly. They are competent. They are operating in good faith. They are short on time. They are not allergic to nuance, but they have limited patience for nuance that does not change what they are going to do on Monday morning. They want to know what the Regulation requires, what it actually expects in practice, where the practitioner consensus has settled, and where the genuinely contested territory still lies. They want defensible answers, not perfect ones, because they have a programme to run and a budget that does not stretch to having every question fully resolved.

If you are a privacy lawyer, you may still find parts of this book

useful, particularly the operator's view of how programmes succeed or fail in practice, but you are not the primary reader, and there will be moments where the framing will feel coarse to you. That trade-off is deliberate. The book that helps everyone helps nobody in particular.

Why this book exists

Most GDPR guidance explains the Regulation; this book closes the gap between knowing it and running a programme that actually satisfies it.

Across a lot of teams figuring GDPR out, a few patterns recur with a regularity that has stopped being surprising.

The first pattern is that competent organisations underestimate the gap between *knowing the Regulation* and *running a programme that satisfies it*. The Regulation is, by the standards of European legislation, reasonably readable. A motivated practitioner can work through the recitals and the operative articles in a long weekend and emerge with a fair grasp of what the Regulation says. They then sit down to apply it to their own organisation and discover that the operative articles are the easy part. The hard part is the dozen judgement calls per processing activity that the articles assume you already know how to make.

The second pattern is that procurement-driven compliance produces brittle programmes. An organisation that built its GDPR posture in response to customer questionnaires tends to have answers that are technically correct, contractually defensible, and operationally fictional. The Records of Processing Activities exists as a spreadsheet that was created for a customer audit two years ago and has not been updated since. The Data Protection Impact Assessment template is downloaded from a popular template library and has never been completed for an actual project. The sub-processor list is current as of the day it was last requested. The first time the organisation faces a real test, a breach, a Data Subject Access Request from a determined former employee, an investigation by a supervisory authority, the gap between the documentation and the reality becomes the dominant fact of the response.

The third pattern is that the operational maturity of a GDPR programme does not correlate with the technical sophistication of the or-

ganisation. A two-person consultancy can run a tighter programme than a thousand-person enterprise. A SaaS company with a deep security culture can have a privacy practice that is two years behind its security practice. A multinational with a global privacy office can have stronger documentation in one region and weaker operational discipline in another. The maturity of the programme depends on whether someone has been paying sustained attention to the operational details, not on the budget or the headcount or the polish of the corporate slide deck.

The fourth pattern, mentioned only briefly here because it recurs throughout the book, is that the consequences of weak programmes are increasingly distributed. The headline news stories tend to focus on the large administrative fines under Article 83: the multi-hundred-million-euro decisions that occasionally land against the largest controllers and that draw most of the press coverage. Those decisions are real, and the fining power under the upper tier is substantial: up to €20 million or 4% of global annual turnover, whichever is higher. But the operational reality for most organisations is broader than the fines. It includes the cost of a contested Data Subject Access Request that escalates into a complaint to the supervisory authority. It includes the procurement-cycle friction of failing a customer's privacy due-diligence questionnaire. It includes the increasing willingness of data subjects to bring private claims for non-material damages under Article 82. And it includes the time cost of a supervisory authority investigation, which is a quieter line item than a fine but, in practice, the one that organisations consistently underestimate.

This book is structured around those patterns. It is, fundamentally, a book about how to run a privacy programme well at whatever scale you are actually operating at. It is not a substitute for the Regulation itself, and it is not a substitute for qualified counsel on specific decisions. It is the operator's companion to both.

Four teams, four maturity tiers

The book uses four recurring organisations as its working examples. Each is introduced by name here, and then they recur throughout the chapters that follow. The four organisations span the spectrum that

GDPR actually catches, from a sole practitioner through to a multinational enterprise, and they are based on composite patterns drawn from organisations worked with, advised, or watched from a respectful distance.

Linden Editorial is Mira's consultancy, the one staring at the German questionnaire at the opening of this chapter. One person. Four to six freelancers used on individual projects. Edinburgh-based, with clients split roughly between the UK, the EU, and the US. Mira's privacy posture is what a lot of small operators' postures look like at the start. A vague awareness that GDPR exists. A stronger awareness that it might apply to her. A creeping suspicion that the answer to *do I need to do something about this* is *yes*, even though she cannot yet say what. Linden Editorial is our window into what minimum-viable compliance actually looks like at the smallest practitioner scale; and into the surprisingly hard question of where the floor sits for an organisation that genuinely cannot afford a full programme.

Northwind Analytics is a Series B business-to-business (B2B) analytics platform headquartered in Austin, Texas. Roughly 180 employees, around 40 of whom are engineers. Two-thirds of their customers are in the US; the other third are in the EU, mostly in Germany, the Netherlands, and France. Northwind has been selling into Europe for around five years without anyone paying serious attention to GDPR as anything other than a procurement requirement. A 2025 enforcement action against a similar US vendor (significant fine, public remediation order, board-level visibility) was the moment their General Counsel realised that the programme he had been running was a procurement-questionnaire-response programme, not an operational compliance programme. Northwind is our window into the gap between *we are compliant on paper* and *we are compliant in practice*, and into the year-long catch-up that closing that gap typically requires. Their General Counsel is Rafael Salinas. Their VP of Security is Brennan Sato. They have an EU-based privacy consultant on retainer named Greta Krüger, who is the outside voice that catches the things Rafael misses.

Heliotrope Innovaties B.V. is a renewable-energy systems integrator based in Utrecht. Roughly 410 employees. They design, install, and maintain commercial-scale solar and battery storage systems for cus-

tomers across the Netherlands, Belgium, Germany, France, and Spain. They have had a DPO since 2018: initially Sofie van der Berg, who has now retired, and currently Pieter de Vries, who took the role in 2023 after a previous career in internal audit. Heliotrope is our window into what GDPR looks like in an organisation where the Regulation has always been the framework rather than a foreign import. Their Records of Processing Activities is documented. Their Data Protection Impact Assessment practice is mature. They have handled three modest breach notifications in the last three years, none of which became enforcement actions. Their lead supervisory authority is the Dutch *Autoriteit Persoonsgegevens* (AP). When we want to show what a working, operationally fluent programme actually looks like (what the daily texture of *being a controller* feels like inside an organisation that has internalised the regime) we will visit Pieter.

Argent Technologies Group is a diversified industrial-technology holding company with offices in seventeen countries, around 21,000 employees, and revenue in the order of €8 billion. They stood up a global privacy programme in 2017 ahead of the May 2018 enforcement date. Their lead supervisory authority is the Irish Data Protection Commission (DPC) under the Article 56 one-stop-shop mechanism, by virtue of their EU headquarters in Dublin. Their Chief Privacy Officer is Roisin Cassidy, who runs a privacy office of fourteen people reporting to the General Counsel. They have deputies in Munich and Singapore, a Privacy Counsel and a Privacy Engineering Lead in Dublin and Madrid respectively, and active engagement with the EDPB through industry trade groups. Argent is our window into what the most mature end of the spectrum looks like: the institutional version of privacy, where the question is no longer *do we have a programme* but *how do we run a programme across seventeen jurisdictions without it fragmenting*.

You will not find every team in every chapter. The point of the four-tier framing is not to deliver four parallel narratives. It is to give you a way to locate yourself. You are probably reading this book because your organisation looks more like one of these four than the others. And if you sit between two of them (a thirty-person EU firm with no DPO, say, bigger than a one-person shop but nothing like a mid-

market compliance function) read yourself as the smaller tier for your obligations and borrow from the next tier up for structure; the chapters flag where that line moves. When the chapters work through a topic, they will pick whichever of the four teams best illustrates the topic at its most instructive maturity tier, and they will tell you what changes if you are at a different tier. By the end of the book, you should have a working sense of which team's posture most resembles yours; and, more importantly, where the next tier up looks like from where you are standing.

A few of the teams are also connected to each other in ways that the chapters will exploit. Northwind is a vendor to Heliotrope; Heliotrope uses Northwind's analytics platform for energy-usage telemetry on its installations. That makes Heliotrope a controller and Northwind a processor for the same body of personal data, governed by a Data Processing Agreement that has been the subject of several rounds of negotiation. Argent partners with Heliotrope on industrial Internet-of-Things integration for large commercial installations, which in some configurations creates joint-controllership obligations under Article 26. And Mira, who we met at the kitchen table, has a side engagement with Heliotrope's marketing department editing English-language communications. That makes her a small-operator processor to a mid-market European controller. The procurement questionnaires of the larger party take the configuration entirely seriously, even when the smaller party is one person and a MacBook. These relationships are not contrivances. They are the kinds of multi-party arrangements that the Regulation is designed to handle, and that working practitioners spend a non-trivial portion of their time mapping out.

What this book promises, and what it does not

What this book does is what a good field guide does. It walks the territory with you, points at the landmarks, names the hazards, tells you which of the available paths the experienced walkers tend to favour, and warns you about the places where the map is genuinely incomplete. The walking is still yours to do.

There is, throughout this book, a quiet thread that runs underneath the practical content. The thread is this: GDPR is, at its philosophical

core, a regulation that takes the position that people have a fundamental interest in the personal data that relates to them, and that organisations which process that data are exercising a delegated trust on their behalf. The Regulation is one expression of that position; it is not the only possible expression, and it will not be the last. The competence this book is trying to help you build has two parts. Part one is the Regulation as it currently stands. Part two is the underlying disposition: the readiness to take the interests of data subjects seriously as a first-order consideration, rather than as a constraint to be minimised. The practitioners who run the best programmes tend to have internalised that disposition. The programmes themselves follow from it.

You do not have to share the disposition to read this book usefully. You can run a defensible GDPR programme purely on the basis that the Regulation is the law and you intend to comply with the law. That is a perfectly respectable starting point and the book will serve you. But it would misrepresent the best practitioners not to say that operational excellence tends to correlate with something more than instrumental compliance, and that the something-more is worth noticing.

A note on a moving landscape

The regulatory landscape moves. The text of the Regulation itself has not been substantially amended since adoption, but the body of interpretive material around it has grown continuously. The European Data Protection Board has issued well over fifty sets of guidelines since 2018, on topics ranging from breach notification under Articles 33 and 34 to the territorial scope of Article 3 to the application of consent under Article 7. National supervisory authorities issue their own guidance, often before the EDPB consolidates a unified position. The Court of Justice of the European Union has decided a string of cases that have reshaped how transfer mechanisms are assessed, how legitimate interests are weighed, and how non-material damages are recoverable under Article 82. The UK ICO operates in parallel and has grown steadily more willing to diverge from EDPB positions, and since February 2026 the divergence is statutory as well as interpretive, which is why several chapters of this book carry a separate UK note. A book of this kind captures the practitioner consensus as of its writing; the consensus con-

tinues to evolve. Treat the citations in this book as a starting point for your own current research, not as the final word, and check the EDPB's guideline register and your competent supervisory authority's recent publications before relying on any specific position in a high-stakes context.

What is coming next

The chapters that follow build outwards from the foundations.

Chapter 2 takes the single most important distinction in the Regulation, the difference between a **controller** and a **processor**, and works through what each role actually entails in practice. Most operators have a vague sense of the distinction, but the vague sense breaks down at the edges, and the edges are where most of the operational disputes happen. We will spend time with Mira, who is discovering that she is a processor under several of her client contracts and a controller for her own client database. We will also spend time with Rafael at Northwind. He is renegotiating the Data Processing Agreement template that has been quietly miscategorising the relationship with a large enterprise customer for the last three years.

From there, the following chapters work through the structural questions that any GDPR programme has to settle. The lawful bases for processing under Article 6. The more restrictive regime for special-category data under Article 9. The rights of data subjects under Articles 13 through 22, and the operational mechanics for honouring them. The records and documentation under Articles 28 and 30 that make the programme legible to its own people and to a supervisory authority. The Chapter V rules that govern moving personal data across international borders, and the vendor and sub-processor relationships through which most of those transfers actually run. Then the working machinery of the modern programme: data mapping in practice, consent interfaces and dark patterns, cookies and the ePrivacy boundary, privacy by design for builders, and the marketing, profiling, and automated-decision rules.

We will look at how the Data Protection Officer role under Articles 37 through 39 actually works in practice. When it is required. What good supervision of a DPO looks like at each maturity tier. We will

spend a chapter on the Data Protection Impact Assessment under Article 35: the most consistently mishandled instrument in the Regulation. And a word on where this volume stops: this is Volume 1 of a three-volume set, *The Law and the Compliance Machine*. The security measures under Article 32, the 72-hour breach regime under Articles 33 and 34, enforcement and fines, and the practical mechanics of dealing with a supervisory authority when an investigation, complaint, or follow-up arrives are the work of Volume 2, *Breach, Enforcement, and Operations*. The templates and worksheets that operationalise both volumes ride in Volume 3, the *Operator's Workbook*.

Volume 2 then closes the set's central question, the one organisations at every maturity tier have to keep answering: how do you tell whether the programme is actually working, and what do you change when it is not?

For now, though, we leave Mira at her kitchen table in Edinburgh, three coffees deep, with the German questionnaire still open on her laptop and the cursor blinking in the field next to *Name and contact details of the Data Protection Officer*. She does not yet know whether she needs one. She does not yet know that the right answer for her is almost certainly *no*, and that she will need to be able to explain why. She does not yet know that the questionnaire is the first of many such questionnaires, and that completing this one well will make every subsequent one easier.

She does know, having read the document for a fourth time, that she is going to have to learn what a controller is, what a processor is, and which one she is for which client. So that is where we are going next.

Recap: GDPR is a directly-applicable EU Regulation in force since May 2018, with a substantively-identical UK counterpart enforced by the ICO. It catches more organisations than people expect, and its operational demands typically outrun its textual demands. This book is for the practitioners who have to run the programme, at any of four maturity tiers, rather than for the lawyers who interpret the Regulation.

Next: Chapter 2: Controller and Processor.

You've just read Chapter 1 of

GDPR Compliance

The full book runs 336 pages across 15 chapters,
with a complete back-of-book index.

Published by Sylvan Press, the book imprint of Sylvan Assurance:
field-tested, plain-English references for the people who do the work.

Scan to see the full book - and where to buy it.
sylvanassurance.com/go/gdpr-compliance

Not ready for the full book? See where you stand first - free, a few minutes:

sylvanassurance.com/go/free-gdpr-checklist

sylvanassurance.com/go/free-gdpr-breach-response