


NOTHING TO HIDE

THE MYTH OF PRIVACY
AND HOW TO RECLAIM IT



TIM GOSKA

TIM GOSKA

NOTHING TO HIDE

THE MYTH OF PRIVACY
AND HOW TO RECLAIM IT

Disclaimer

Cyber Konekt has affiliate relationships with some products mentioned in this book, including NordVPN and Surfshark. If you purchase through certain links, Cyber Konekt may receive a commission at no additional cost to you. These relationships may create a commercial interest in recommendations. The inclusion, ranking or discussion of any product is based on the criteria stated in this book, and readers should independently assess whether a product is suitable for their needs.

Product comparisons in this book are for educational purposes and are based on publicly available information, provider documentation, published research, regulatory findings and/or testing available at the time of writing. Services, prices, privacy policies and technical features may change.

All characters, case studies and scenarios in this book are fictionalised composites created solely for illustrative and educational purposes. They are not intended to depict, identify or refer to any specific real person, living or deceased. Any resemblance to an actual person is coincidental. The scenarios are based on general patterns, publicly documented practices, reported case studies and common privacy-risk situations, with identifying details altered or omitted.

This book is provided for general educational and informational purposes only. It is not legal, regulatory, cybersecurity, financial or professional advice tailored to any individual, organisation or jurisdiction.

Laws, regulatory guidance, platform practices and product features change frequently. The information reflects the author's understanding at the time of writing. Readers should verify current requirements and obtain appropriate professional advice before acting on any information in this book.

Nothing to Hide

The Myth of Privacy and How to Reclaim It

Copyright © 2026 Tim Goska
All rights reserved.

This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International Licence (CC BY-NC 4.0). You are free to share and adapt this material for non-commercial purposes, provided you give appropriate credit to the author and indicate if changes were made. Adapted versions are not endorsed by Cyber Konekt unless expressly approved. For commercial use, please contact cyberkonekt.com. Full licence terms at creativecommons.org/licenses/by-nc/4.0.

First published 2026
Czech Republic

ISBN 978-80-11-08922-1 (paperback)

ISBN 978-80-11-08923-8 (pdf)

ISBN 978-80-11-08924-5 (ePub)

Published by Cyber Konekt | www.cyberkonekt.com

All trademarks, service marks, and company names mentioned in this book are the property of their respective owners. Their inclusion does not imply endorsement of this book or its contents.

For permissions requests and rights enquiries, please contact:
support@cyberkonekt.com

Contents

Introduction	10
--------------	----

Part 1 The Problem

1	Your Digital Shadow	16
2	The Business of Watching	26
3	It Is Worse in Some Places Than Others	37
4	Real People, Real Harm	49

Part 2 How It Got This Bad

5	Convenience Was Always the Trojan Horse	60
6	The Fingerprinting Arms Race	70
7	Why the Rules Are Not Working	79

Part 3

What You Can Do

8	The Layered Defense	91
9	Your Files, Your Communications	102
10	Different Threats for Different Lives	113
11	How to Spot a Privacy Tool That Is Lying to You	127

Part 4

What Comes Next

12	What Good Regulation Looks Like	139
13	The Companies Building Differently	152
14	You Are Not Powerless	164

A Note From the Author	173
Notes & Sources	177
Glossary	198

Introduction

As of 2026, approximately 6.64 billion people, roughly 82 percent of the global population, are covered by at least one data protection law. That is more legal protection for personal privacy than at any point in human history.

And yet the surveillance economy has never been larger. GDPR fines alone have exceeded 7.1 billion euros across all EU member states since 2018, and approximately 40 percent of that total is currently annulled or under active legal challenge, which tells you both that violations are widespread and that the consequences remain insufficient to stop them. Over 80 percent of free mobile apps track private user data. Free apps are up to four times more likely to collect your data than paid ones. The average person's phone generates thousands of data points every day, most of them invisible, none of them meaningfully consented to.

More protection on paper. More surveillance in practice. That tension is what this book is about. But before we get into how the system works, I want to address the thought that stops most people from engaging with this subject at all.

I have nothing to hide. So why does it matter?

It is a reasonable thing to think. Most people are not doing anything illegal. Most people do not feel watched. The services are genuinely useful, often excellent, and the trade feels abstract: some data you cannot see, in exchange for tools you

use every day. If you are not hiding anything, what exactly is the problem?

I want to take that question seriously, because it deserves a serious answer rather than a dismissal.

The first assumption is that privacy is about concealment.

It is not. Privacy is about control. You close the bathroom door not because you are doing anything wrong, but because some spaces are yours to manage. Nobody interprets a closed bathroom door as evidence of criminal activity. The same logic applies to your search history, even if the law has not quite caught up to that. You do not share your salary with everyone you meet, not because you are ashamed of it, but because that information has different consequences depending on who holds it. Privacy is the right to decide who knows what about you, and when, and for what purpose. The "nothing to hide" framing collapses that into a binary: either you have something to hide, or you have nothing to fear. But that is not how information works, and it is not how power works.

The second assumption is that information stays in the context where you created it. It does not. Your doctor knowing you have anxiety is appropriate. Your employer knowing it is a different matter. Your insurer knowing it changes what you pay. Your landlord knowing it may affect whether you are approved for a lease. The data you generate online does not stay where you left it. It moves, is combined with other data, is inferred from, and appears in contexts you never anticipated and would never have chosen. You do not have "nothing to hide" in some pure abstract sense. You have

information that is appropriate in some contexts and damaging in others, and the surveillance economy systematically strips those contextual boundaries away. What feels like a harmless search today becomes part of a profile that makes decisions about you tomorrow, in rooms you will never enter, by people you will never meet.

The third assumption is that you are consenting for the person you are right now. You are not. Data is not perishable. The behavioral record being built about you today will exist in whatever future you inhabit, under whatever laws, administered by whatever government, in whatever political climate prevails ten, twenty, or thirty years from now. What is legal today may not be tomorrow. What is politically neutral today may be dangerous under a different administration. What feels like an innocuous preference profile in a stable democracy looks entirely different in a country moving in another direction. You are not just saying "I have nothing to hide today." You are making a permanent decision on behalf of every version of yourself that will ever exist, in every future you cannot predict. History is consistent on this point: data collected under one political regime has a habit of becoming very useful under the next one.

The fourth assumption is that the exchange is fair. It is not. "Nothing to hide" treats the transaction as neutral: you give up some information, you get a service, no one is harmed. But the party on the other side of that exchange has an enormous amount of information about you. You have almost none about them: what they hold, who they have shared it with, how long they keep it, what they infer from it, who has bought it, or

what decisions it has already influenced about you. You are negotiating in the dark, in a system designed by people with complete information, operating under incentives that are not aligned with yours. A contract signed without knowledge of its terms is not a fair contract. The "nothing to hide" framing assumes a fair exchange. The exchange is not fair.

Most people who care about this issue have not acted on it. Not because they do not want to. Because the problem feels too large, too technical, and too entrenched to approach from the outside. This book is written for exactly those people. It assumes no technical knowledge. It does not require you to already be convinced that privacy matters. In fact, if you have just read the four points above and still feel uncertain, good. Uncertainty is a more honest starting position than either panic or indifference, and this book is designed to work from exactly that place.

I am the co-founder and CTO of Cyber Konekt, a privacy-first technology company based in Prague. I have spent years building tools designed to protect the people who use them, making the decisions that a genuine privacy commitment actually requires, in a market where the surveillance-funded free alternative is always one click away. The perspective in this book is not that of an observer. It is that of someone who has built in this space and faced the real choices that building here involves.

The surveillance economy is large, technically sophisticated, and economically powerful. It is not invincible. It depends on your passivity. It depends on you not knowing how it works. It

depends on the assumption that you will not bother.

Sara is waiting at the tram stop. Let us begin.

Tim Goska

Co-founder and CTO, Cyber Konekt

Prague, Czech Republic

2026

PART 1

THE PROBLEM

CHAPTER 1

YOUR DIGITAL SHADOW

Sara leaves her apartment in Prague at 7:43 on a Tuesday morning. She has not opened a single app. She has not searched for anything. She has not typed a single word into any device. She locks the door, puts her phone in her coat pocket, and walks to the tram stop.

By the time she sits down in the tram, fourteen separate data points about her have already been collected.

Her phone has pinged three cell towers, giving her precise location to within a few metres, logging the time of each ping, and identifying her network provider. That is five data points before she has taken ten steps. Her building's smart entry system has logged the exact time she left and the unique identifier of her key fob. A Bluetooth sensor installed by an advertising network inside a nearby shop window, has detected the unique identifier broadcast by her phone and recorded how long she paused nearby. A traffic camera operated by the city has captured her face and her direction of travel. And the tram's ticketing system, which she paid with a contactless bank card, now knows her route, her direction, the time she boarded, and the last four digits of the card she used.

Sara has done nothing unusual, she is not being investigated, nobody has issued a warrant or filed a request. This is just a Tuesday.

This is just how it works now.

The Shadow You Did Not Know You Were Casting

Most people, when they think about privacy, think about passwords. They think about not sharing their home address with strangers, or being careful about what they post online, these instincts are reasonable. They are also, at this point, almost beside the point.

The data being collected about you today, is not the data you consciously hand over. It is the data generated by the simple act of existing in a world full of connected devices, networked sensors, and companies whose entire business model depends on knowing as much about you as possible.

Your phone knows where you sleep, because that is where it spends the night. It knows where you work, because that is where it spends most weekday mornings. It knows who you spend time with, because it can see which other phones are consistently nearby. It knows, with surprising accuracy, whether you are anxious or calm, rushing or relaxed, based on how you move and how fast you type. None of this requires you to tell it anything. It simply watches.

This is what I mean when I talk about your digital shadow. It is not the identity you present online. It is the one you create without trying. It follows you everywhere, grows longer as the day goes on, and is being read by people and systems you will never meet.

The good news, and I want to say this clearly at the start of this book, is that you can shrink it. Not eliminate it entirely, because the world we live in makes that nearly impossible without significant sacrifice. But shrink it, shape it, and take back meaningful control over it. That is what this book is about.

What Is Actually Being Collected

Let us be specific, because vague warnings about 'your data' are part of why most people tune out. When I say data is being collected about you, I mean something precise.

Location data is perhaps the most revealing. A study published by researchers at the MIT Media Lab found that a dataset of location records, with names removed, could be re-identified with over 95 percent accuracy using nothing more than four data points. Your home, your workplace, a shop you visit regularly, and a gym or school or place of worship. Four places and you are identifiable. Most people's phones generate thousands of location records every day.

Behavioral data is collected by almost every app you use. How long you spend on each screen. Which articles you read to the end and which you abandon halfway through. What time you wake up, based on when you first unlock your phone. Whether you scroll quickly, suggesting anxiety or distraction, or slowly, suggesting focus. These signals are not collected because any single person cares about your scrolling speed. They are collected because, in aggregate, they are extraordinarily good at predicting what you will do next. And predicting behavior is worth an enormous amount of money.

Device fingerprinting is one that most people have never heard of, and it is worth spending a moment on. Every device you own has a slightly unique combination of hardware, software settings, screen resolution, installed fonts, and dozens of other technical characteristics. Together, these form a fingerprint that can identify your device, and therefore you, even if you clear your cookies, use a private browsing window, or switch to a different browser. I have spent years building tools to fight exactly this, and I can tell you: it is more persistent than most people realise.

And then there is the data you do hand over willingly: the forms you fill in, the accounts you create, the loyalty cards you swipe, the surveys you answer for a ten percent discount. This data is legal, consensual in a technical sense, and traded between companies in ways that would take a legal team several hours to map out from a single privacy policy. Most people never read those policies. Most people have no idea that the supermarket chain they use sold their purchase history to a health insurance data broker. Most people, frankly, have no reason to suspect this is happening.

But it is.

The Map of Who Is Watching

It helps to understand that data collection is not one thing done by one type of organisation. It is a layered system with several distinct players, and they often do not even know about each other.

At the top you have the platforms: the big technology companies whose products you use every day. Search engines, social networks, email providers, map services, video platforms. These companies collect data directly from your use of their products, and their scale is extraordinary. One major search engine processes over 3.5 billion searches every day, amounting to more than 1.2 trillion searches per year worldwide. Each one of those queries tells that company something about the person who typed it.

Below them, and far less visible, sit the data brokers. These are companies most people have never heard of, operating in a market worth over 300 billion euros globally, according to industry analysts. They do not make products you use, their product is you. They buy data from hundreds of sources, clean it, combine it, and sell detailed profiles to anyone willing to pay: insurers, employers, advertisers, political campaigns, debt collectors, and, in countries with weaker regulations, governments.

Then there are the device manufacturers, the operating system providers, the app developers, and the network operators. Each layer collects something. Sometimes they share it. Sometimes they sell it. Sometimes they simply hold it, waiting for a use case they have not thought of yet.

The result is a surveillance infrastructure that no single person designed, and no single organisation controls. It emerged because data is valuable, and collecting it is cheap, and for a long time there were almost no rules about what you could do with it.

Some of those rules now exist. The GDPR, which came into force across the European Union in 2018, was a genuine shift. It gave people real rights over their data and imposed real penalties for violations. But it is imperfect, unevenly enforced, and only covers one region. In the United States, there is still no single federal privacy law. In many parts of Southeast Asia and Africa, comprehensive frameworks are only beginning to emerge. Globally, the picture is a patchwork of protection and exposure, and where you happen to live shapes your privacy more than almost any personal choice you can make.

Why This Matters More Than You Might Think

I want to be honest with you about something. When I started building privacy tools, I had a fairly technical view of this problem. I cared about the architecture. I cared about the clever ways tracking systems worked and the clever ways they could be defeated. It took me longer than I would like to admit to fully understand the human dimension.

Privacy is not really about data. It is about power.

When someone knows more about you than you know about them, the relationship is unequal. When a company knows your financial anxieties, your health concerns, your relationship problems, and your daily routine, it can use that knowledge to manipulate your decisions in ways you cannot see or counter. When a government has access to the same information, the implications are more serious still. History is not short of examples of what states do with detailed records of their citizens' lives, beliefs, and associations.

This is not an argument that everyone collecting your data has bad intentions. Many do not. But intentions change, companies get acquired, databases get breached, and the data that felt harmless when it was collected can become something quite different in different hands or a different political climate. The researcher Shoshana Zuboff, who spent years studying this ecosystem, described it as 'the raw material of a new economic order.' She was right. The question is who that order serves.

There is also something quieter and harder to name. When you know you are being watched, you behave differently. Studies of online behaviour consistently show that awareness of surveillance changes what people search for, what they say, and which ideas they are willing to explore. Surveillance does not just record who you are. Over time, it shapes who you become.

The Two Minutes That Actually Help

Before we go any further, I want to give you one thing you can do right now, because I think the best way to earn your trust as a reader is to be useful immediately.

Go to your phone's settings and find the section for location permissions. On most phones this is under Privacy or Apps. Look at how many applications have been granted 'always on' location access, meaning they track you even when the app is not open. For most people, the number is somewhere between unsettling and genuinely alarming. Revoke that access for any app that does not have an obvious need for your location at all times. A weather app needs your location when you open it.

It does not need it at 3am.

That one change will not solve everything. But it will reduce, in a concrete and measurable way, the amount of data flowing out of your device without your awareness. It takes about two minutes, and it will probably change how you think about the other permissions sitting quietly on your phone.

What the Rest of This Book Does

Chapter 1 is the hardest chapter to write, because it has to do two contradictory things at once. It has to show you that the problem is serious enough to pay attention to. And it has to leave you feeling that doing something about it is actually possible.

I believe both of those things are true. I have spent years working on this, and I have seen what good tools, good habits, and better regulation can do. The surveillance infrastructure is vast, but it is not invincible. It depends on your passivity. It depends on you not knowing how it works. It depends on the assumption that you will not bother.

The rest of this book is about bothering.

Part One continues with how the business of watching actually works, and who is profiting from it. Part Two explains how the current system came to exist and why fixing it has proven so difficult. Part Three gives you a practical, non-technical toolkit for taking back control. And Part Four looks at what is changing, what is possible, and what you can push for beyond your own devices.

Sara, on her Tuesday tram to work, does not know that fourteen data points about her morning have already been collected. But she could know. And knowing, it turns out, is most of the battle.

Five Things You Can Do This Week

1. Audit your location permissions. Revoke 'always on' access for any app that does not genuinely need it.
2. Check which apps have access to your microphone and camera. Revoke anything you do not recognise or no longer use.
3. Search your own name in three different search engines and see what appears. This is roughly what a data broker profile of you looks like from the outside.
4. Read the first paragraph of any privacy policy you have clicked 'agree' on recently. Just the first paragraph. You will learn something.
5. Switch your default search engine to one that does not track your queries. It takes about thirty seconds and the search results are, genuinely, fine.

The digital world you just woke up to has been watching you since before you opened this page. It will still be watching after you close it. The difference is that now you know.

CHAPTER 2

THE BUSINESS OF WATCHING

Imagine you are sitting in a meeting room on the forty-second floor of a glass office building in San Francisco. It is 2012. Around the table are a handful of engineers and a few people from the business side, and someone is presenting a slide that shows a graph going up and to the right. The graph represents something called the daily active user count, and it is climbing fast.

Someone asks the obvious question. How do we make money from this?

The answer, it turns out, is both simple and world-altering. You make money by learning everything you possibly can about the people represented by that climbing line, and then you sell access to their attention to the highest bidder. You do not charge users a subscription. You do not sell a product. The users are the product. Or more precisely, what you know about them is the product, and advertisers are the actual customers. That meeting, or one very much like it, happened at dozens of companies across Silicon Valley in the early 2010s. It was not a conspiracy. Nobody sat in a dimly lit room and decided to build a global surveillance apparatus. It emerged, incrementally and almost accidentally, from a set of very rational short-term decisions made by people who were mostly just trying to build something people wanted to use.

The fact that it was not a conspiracy makes it, in some ways, harder to fight.

How the Model Actually Works

Surveillance capitalism, a term coined by the Harvard professor Shoshana Zuboff, is not really about surveillance in the dramatic sense. There are no agents following you. Nobody is listening to your phone calls, at least not in the commercial context we are discussing here. What it describes is something more mundane and, in its own way, more insidious: the systematic extraction of behavioral data from human experience, and the conversion of that data into predictions about what you will do next.

Predictions are the product. Not the data itself, which is merely the raw material. What advertisers pay for is the ability to reach you at a specific moment, with a specific message, when you are statistically most likely to respond to it. A company selling running shoes does not want to show you an advertisement when you are looking at photographs of your cousin's wedding. They want to show it to you thirty seconds after you have searched for a half-marathon training plan, or two minutes after your fitness tracker recorded your morning run.

To build that kind of precision, you need a lot of data. You need to know not just what people search for, but how long they pause on certain pages. Not just what they buy, but what they almost bought and then did not. Not just where they go, but the routes they take and the times they travel and whether they deviate from their usual patterns. The system is not trying to understand you as a person. It is trying to build a model of your behavior that is accurate enough to be commercially useful.

This is an important distinction. Nobody at these companies is reading your messages. The scale makes that impossible and, frankly, unnecessary. What the system produces is something closer to a very detailed weather forecast for your behavior. Seventy percent chance this user clicks on a discount offer between 6pm and 8pm on a weekday. High probability of a purchase if shown social proof from peers in the same income bracket. Elevated anxiety score today, based on search and scrolling patterns, suggesting receptivity to certain categories of financial products.

The forecast does not need to be perfect. It just needs to be better than guessing.

Who Is in the Room

The surveillance economy is not one industry. It is an ecosystem, and understanding its different layers is part of understanding why it is so difficult to opt out of.

The platforms are the most visible layer. These are the companies whose products you actually use: search engines, social networks, email services, map applications, video platforms, and the operating systems running on your devices. They collect data directly from your interactions with their products, and their reach is extraordinary. One major social platform has more than 3.5 billion daily active users across its family of apps. To put that in perspective: that is more people than live in China and India combined. When a single company has behavioral data on that portion of the human population, the phrase 'market power' starts to feel inadequate.

Below the platforms, and largely invisible to ordinary users, sits the data broker industry. These companies do not make anything you use. Their entire business is the collection, cleaning, combination, and resale of personal data. A report from the Norwegian Consumer Council in 2020 analysed ten popular Android apps and found they collectively transmitted user data to at least 135 different third parties within the advertising space, covering 216 unique domains.

Then there are the advertising technology companies, known as ad tech, which operate the infrastructure connecting the platforms and the data brokers to the advertisers. This is the plumbing of the surveillance economy, and it is almost entirely hidden from public view. When you load a webpage and see an advertisement, a process called real-time bidding has typically occurred in the milliseconds between you clicking the link and the page appearing on your screen. Your profile has been assessed, your predicted value to multiple advertisers has been calculated, and the highest bidder has won the right to show you something. All of this happens before you have finished reading the headline.

It is, from a purely engineering perspective, an astonishing piece of infrastructure. I say that without admiration.

The Number That Should Concern You

The global data broker market was valued at over 300 billion euros globally according to industry research. It is projected to roughly double by 2030.

To understand what that money represents, consider what is being sold. Data brokers hold files on most adults in the developed world. A typical file might include your name and address history, your estimated income and net worth, your educational background, your purchasing history across dozens of retailers, your political affiliation inferred from your reading habits, your health conditions inferred from your search history, your relationship status, your religious views, and scores of other attributes derived from the data you have generated over years of living a connected life.

Most people have never consented to this in any meaningful sense. They clicked through a terms and conditions document they did not read, on a platform they wanted to use, which sold their data to a partner they had never heard of, which sold it on again. The consent is technically present in the paperwork. Whether it is real consent by any reasonable definition is a different question.

The regulations are beginning to catch up. Under the GDPR in Europe, citizens have the right to request access to the data held about them, the right to have it corrected, and in certain circumstances the right to have it deleted. Several data brokers have been fined for non-compliance. In California, the Consumer Privacy Act gives residents similar rights.

In South Korea, the Personal Information Protection Act is among the strictest in the world. Progress is real.

But regulation moves slowly, and the data economy moves fast. The rules written in 2018 were not written for the capabilities that exist in 2026. And in most of the world, comprehensive privacy protection simply does not exist.

The Part Nobody Talks About Enough

There is a dimension of this that gets less attention than it deserves, and I want to spend a moment on it.

When we talk about data collection and targeted advertising, it is easy to frame it as a mildly annoying but essentially harmless exchange. You get free services. Companies get to show you relevant ads. Nobody gets hurt.

But the same infrastructure that delivers you a running shoe advertisement at exactly the right psychological moment can also be used to deliver a political message designed to inflame your existing fears, a financial product designed to exploit your anxiety, or health misinformation designed to look like the content you already trust. The targeting capability does not know, or care, whether what it is delivering is useful to you or harmful to you. It only knows whether you engaged with it.

Engagement, in the language of the attention economy, means clicks, shares, time spent, and emotional responses measured through behavior. Anger and fear drive more engagement than calm and contentment. This is not a design choice anyone made deliberately. It is an emergent property of optimising for

attention without any constraint on what kind of attention counts.

The Frances Haugen leak of internal Facebook documents in 2021 included research showing that the company's own teams had identified that its algorithms were amplifying divisive content because it drove higher engagement, and that this effect was contributing to real-world harm in multiple countries. Haugen testified and leaked internal documents, showing that the company was aware the engagement-based system was causing real-world harm, and that promoting safety would reduce the time users spent on the platform and therefore reduce ad revenue. The company continued prioritising engagement anyway.

I am not trying to make you angry at any particular company. Anger is not especially useful here. What I want you to understand is that the surveillance economy is not a neutral infrastructure that sometimes gets misused. It is a system that, by its nature, trades in human vulnerability. That is not a side effect. It is the mechanism.

Why Free Is the Most Expensive Price

The metaphor I want you to carry through this book is this: the internet you have been using is not a public service. It is a shopping centre.

When you walk into a shopping centre, everything is designed to keep you there as long as possible and to make you spend as much as possible. The layout is not accidental. The lighting is chosen carefully.

The smells near the food court are not random. Every element of the environment has been optimised, based on decades of behavioral research, to influence your decisions without your awareness.

Now imagine that shopping centre also had cameras in every corner, tracking not just where you walked but how long you paused at each window, which displays made your pupils dilate, and which other shoppers you seemed to know. Imagine it shared that footage with three hundred other businesses. Imagine you could never leave because every other place you went was somehow connected to the same network.

That is roughly the digital environment most of us inhabit. And we accepted it because the shops were free to enter, the entertainment was free to watch, and the maps were free to use. Free, in this context, turned out to be the most expensive price of all.

The good news, and I mean this, is that the shopping centre has fire exits. They are not well signposted. Some of them are a bit inconvenient. But they exist, and they work.

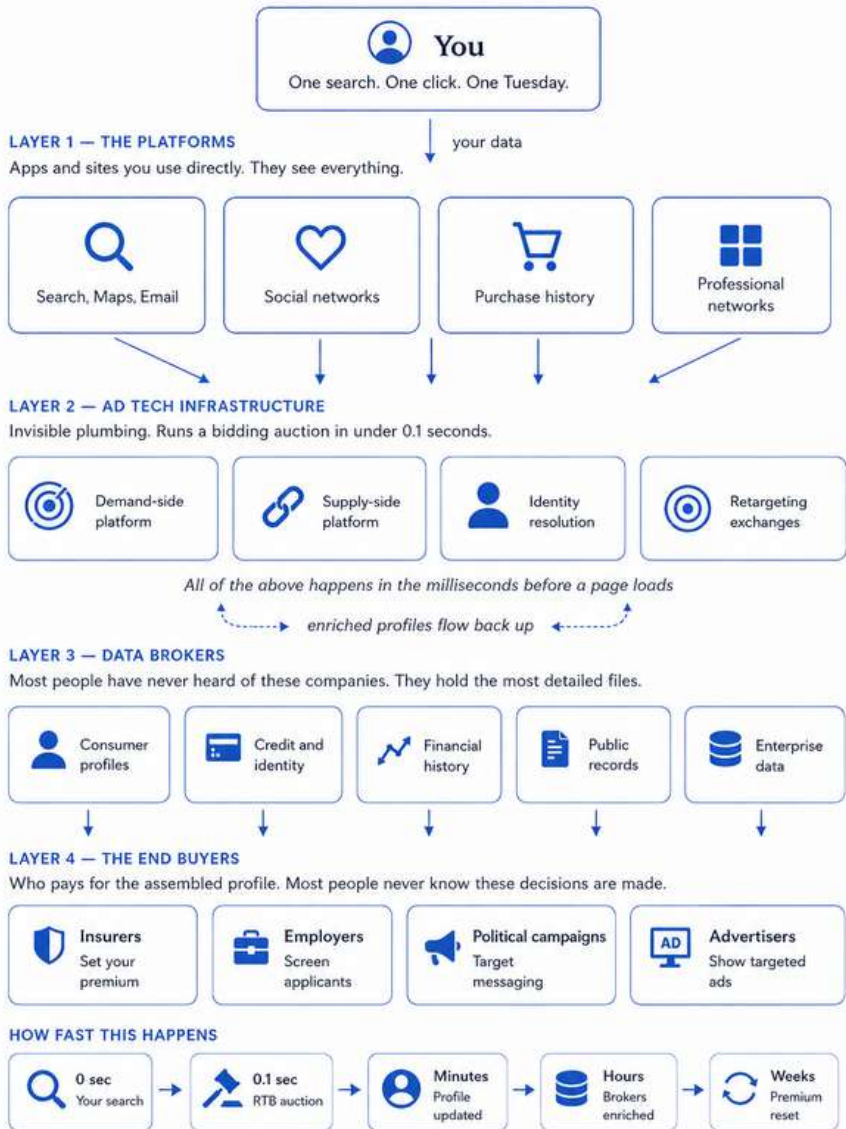
We've now mapped the terrain: data brokers stockpile personal data, ad-tech engines turn that data into revenue, and regulators are scrambling to put caps on the practice. But what does that mean for you? The rest of this book will explore how to locate the 'fire exits' in the data-packed shopping centre, and how to actually use them.

Five Things Worth Knowing

1. You can request a copy of the data a broker holds on you. In Europe, GDPR gives you this right. In the US, California residents can do this under the CPRA. It is worth doing once, just to see.
2. Real-time bidding, the system that auctions your attention in milliseconds, is currently under legal challenge in both the EU and the UK on the grounds that it violates data protection law by design.
3. Most data brokers allow opt-out requests, though the process is deliberately cumbersome. Services exist that will submit these requests on your behalf automatically.
4. Your phone's advertising ID, a unique identifier used to track you across apps, can be reset or disabled entirely in your device settings. It takes about thirty seconds.
5. The business model described in this chapter applies to free services. Paid alternatives exist for search, email, and cloud storage that are funded by subscriptions rather than data. They are not perfect, but the incentive structure is different.

Sara's data, sold to four companies before she reached her office, will outlast the phone she carried it on.

The Surveillance Infrastructure



The United States alone houses over 4,000 active data brokerage entities. The top ten collectively account for approximately 41% of global revenues.

CHAPTER 3

IT IS WORSE IN SOME PLACES THAN OTHERS

In May 2018, a French lawyer named Isabelle sat at her desk in Lyon and did something that, until that week, would have been essentially pointless. She wrote a formal letter to a large retail company asking them to tell her exactly what personal data they held about her, where they had obtained it, who they had shared it with, and on what legal basis they were processing it. She received a detailed written response within three weeks. It listed every category of data they held, the third parties it had been shared with, and the retention period for each type. At the bottom, it reminded her that she had the right to request deletion of any of it.

Isabelle had these rights because, on 25 May 2018, the General Data Protection Regulation came into force across the European Union. The GDPR, as it is almost universally known, was four years in the making and represents the most comprehensive attempt by any government body in the world to give citizens genuine legal control over their personal data. It is not perfect. But it is, by a considerable distance, the closest thing the world has to a gold standard.

The reason I am starting with Isabelle in Lyon rather than somewhere with weaker protections is deliberate. I want you to understand what is actually possible before we look at how far most of the world falls short of it. Because once you know what good looks like, the gaps become much harder to accept.

What the GDPR Actually Does

The GDPR is a long and technical document, and you do not need to read it. But understanding its core principles takes about five minutes and is genuinely worth the time, because these principles are the framework against which every other privacy regime in the world can be measured.

The first principle is lawful basis. Under the GDPR, a company cannot simply collect your data because it wants to. It must have a specific legal reason for doing so. The most common basis is consent, which must be freely given, specific, informed, and unambiguous. A pre-ticked box does not count. A statement buried in paragraph forty-seven of a terms and conditions document does not count. You must be told clearly what you are agreeing to, and you must actively agree to it.

The second principle is data minimisation. Companies can only collect the data they actually need for the stated purpose. They cannot vacuum up everything available on the grounds that it might be useful later. This one principle, if consistently enforced, would fundamentally change the economics of surveillance capitalism.

The third is individual rights. Under the GDPR, every person in the EU has the right to access their data, correct it, restrict its processing, and in certain circumstances delete it entirely. The right to erasure, sometimes called the right to be forgotten, made headlines when it was introduced. It is more limited in practice than the name suggests, but it is real and it has been used.

And the fourth is enforcement with teeth. Fines under the GDPR can reach four percent of a company's global annual turnover, or twenty million euros, whichever is higher. In 2023, Meta was fined 1.2 billion euros by the Irish Data Protection Commission for transferring European user data to the United States without adequate safeguards. That is not a rounding error on a quarterly earnings report. That is a number that changes board-level conversations.

The GDPR has problems. Enforcement is inconsistent across member states, with some national regulators moving far more slowly than others. Smaller companies often lack the resources to achieve full compliance. And the consent mechanisms that the regulation intended to be meaningful have, in many cases, been degraded into the cookie banners you now dismiss without reading on every website you visit. The mechanism was right. The implementation has been imperfect.

But the direction of travel is correct. And the GDPR has had an effect beyond Europe's borders that its authors probably hoped for but could not have guaranteed.

The Brussels Effect

When the GDPR came into force, a quiet and largely unreported thing happened. Companies that operated globally found themselves facing a choice: maintain one set of data practices for European users and a different set for everyone else, or simply apply the European standard everywhere.

Many chose the latter. Not out of altruism, but because maintaining two parallel compliance systems is expensive and

complicated, and because the reputational risk of being seen to treat non-European users as second-class was growing. The result was that the GDPR's protections quietly extended to millions of people who had no legal right to them.

This phenomenon has a name in regulatory circles: the Brussels Effect. The political scientist Anu Bradford coined the term to describe the EU's tendency to set global standards through market power rather than formal international agreements. The EU is the world's largest single market. If you want to operate in it, you follow its rules. And if you are already following its rules, you often end up following them everywhere.

The Brussels Effect is one reason privacy advocates have been cautiously optimistic about the GDPR's long-term influence. But it is not a substitute for law. It is a market dynamic, and market dynamics can reverse. What one company chooses to apply voluntarily, another company can choose not to.

The United States: A Patchwork with Holes

The United States has no federal privacy law. This is not an oversight. It is the outcome of decades of lobbying by technology and advertising companies, combined with a political culture that has historically been more suspicious of government regulation than of corporate data collection.

What exists instead is a patchwork of sector-specific federal laws and a growing collection of state-level regulations. HIPAA covers health information. COPPA covers data collected from children under thirteen. FERPA covers educational records.

The Gramm-Leach-Bliley Act covers certain financial data. Each of these laws protects a specific category of information in a specific context, and the gaps between them are wide enough to drive a data broker's server farm through.

California has gone furthest of any US state. The California Consumer Privacy Act, strengthened in 2023 by the California Privacy Rights Act, gives Californians rights that look broadly similar to GDPR provisions: the right to know what is collected, the right to delete, the right to opt out of the sale of their data. Several other states, including Virginia, Colorado, Texas, and Connecticut, have passed their own versions. But these laws vary significantly in scope, in the rights they grant, and in how vigorously they are enforced.

The practical result is that your privacy rights in the United States depend heavily on which state you live in, which sector the company collecting your data operates in, and whether that company has been unlucky enough to attract the attention of a regulator. That is not a privacy framework. It is a lottery.

Federal legislation has been proposed repeatedly over the past decade. It has not passed. The American Data Privacy and Protection Act came closest in 2022, clearing a committee with rare bipartisan support before stalling in the full Congress. The reasons it failed are instructive: disagreement over whether a federal law should override state laws like California's, and sustained resistance from technology and advertising industry groups who argued, as they always do, that regulation would harm innovation.

The Rest of the World: Progress, Gaps, and Danger Zones

Beyond Europe and the United States, the picture becomes significantly more varied, and in places, significantly more concerning.

Several countries have built strong frameworks explicitly modelled on the GDPR. Canada's updated privacy legislation, currently working through Parliament under Bill C-27, significantly strengthens rights under the existing PIPEDA framework. Brazil's Lei Geral de Proteção de Dados, known as the LGPD, came into force in 2020 and closely mirrors GDPR principles. South Korea's Personal Information Protection Act is considered one of the strictest in the world. Japan, Argentina, and New Zealand all have frameworks that provide meaningful protection, though each has gaps.

India passed its Digital Personal Data Protection Act in 2023, covering a population of 1.4 billion people. The law grants individuals meaningful rights over their data and imposes obligations on companies processing it. Privacy advocates have raised concerns about the broad exemptions granted to the government, which can override protections in the interests of national security. This tension between individual privacy and state access runs through almost every privacy framework in the world, but it is more pronounced in some places than others.

In much of sub-Saharan Africa, Southeast Asia, and the Middle East, comprehensive privacy legislation either does not exist or exists only on paper. This does not mean that people in these regions do not care about their privacy, or that their data is not valuable. It means that the companies and governments collecting their data face fewer legal consequences for misusing it. The global data economy does not pause at borders. The protections do.

And then there are the countries where the absence of privacy protection is not an oversight but a policy. China's data ecosystem is vast, technically sophisticated, and oriented primarily toward state control. The Skynet surveillance network has grown from 20 million cameras in 2018 to a planned 570 million, blanketing public spaces with facial recognition infrastructure capable of identifying individuals in real time. The Sharp Eyes program extends this reach into rural communities, turning neighbours into surveillance participants through cameras connected to residential television sets. In Xinjiang, the Integrated Joint Operations Platform, first exposed by Human Rights Watch in 2019 and significantly upgraded since, monitors the movements, communications, financial transactions, and social connections of Uyghur and other Muslim minority populations. Former Xinjiang government engineers have described a system that flagged thousands of people for detention based on algorithmic assessment alone, without evidence of any crime. This is not a hypothetical surveillance state. It is an operational one, and it is causing documented, serious harm to real people right now.

The Social Credit System, which tends to dominate Western coverage of Chinese surveillance, is a different and more complicated picture. It is worth understanding accurately, not because the concerns are overblown, but because the imprecision in how it is typically described obscures both what is genuinely alarming and what is not. What actually exists is a fragmented collection of local and sectoral pilots, primarily focused on corporate compliance, financial defaulters, and court order violators. A unified citizen score tracking all individual behaviour nationally has never been implemented. As of early 2026, most local individual scoring pilots have ended. Jeremy Daum, Senior Fellow at Yale Law School's Paul Tsai China Center and the leading English-language authority on this system, describes it as closer to an administrative enforcement infrastructure than a surveillance panopticon.

The conflation of these two things, the operational surveillance infrastructure and the social credit administrative system, matters precisely because it makes the actual threat harder to see clearly. China's surveillance state does not primarily operate through a citizen score. It operates through cameras, algorithms, databases, and a legal framework that gives the state essentially unlimited access to the data these systems generate. That infrastructure is real, it is growing, and it is being exported. Chinese surveillance technology companies, several of which are now subject to US sanctions for their role in Xinjiang, have sold systems to governments across Africa, Central Asia, and the Middle East. The tools, the techniques, and the architecture of a surveillance state do not stay inside the borders of the country that built them.

China is the most technologically advanced example but it is not the only one. Russia has moved toward data localisation requirements that are less about protecting citizens and more about ensuring the state can access data held within its borders.

I am not raising these examples to be alarmist. Most readers of this book are not living under active state surveillance. But these cases matter for two reasons. First, as highlighted, the infrastructure being built in authoritarian contexts does not stay there. The tools, the techniques, and the companies that build them operate globally. Second, because the direction of travel in democratic countries, if left unchecked, points toward more data collection, more state access, and less individual control. History suggests that capabilities, once created, tend to be used.

Why Where You Live Is a Privacy Decision You Did Not Get to Make

I have spent a lot of time thinking about this asymmetry, because it sits at the heart of why building privacy tools feels urgent to me. The protections you have are largely determined by an accident of geography. A user in Berlin and a user in Jakarta can be customers of the same app, using the same features, generating the same data. One of them has enforceable legal rights over that data. The other has almost none.

This is not the fault of the user in Jakarta. And it is not something that individual action can fully fix. Better tools help. Good habits help. But the structural gap between strong and weak regulatory environments can only be closed by regulation, by international pressure, and by companies choosing to apply consistent standards regardless of what local law requires.

The good news is that the regulatory tide is moving in one direction globally, even if it is moving slowly and unevenly. The number of countries with comprehensive privacy legislation has grown from around fifty in 2010 to more than 140 countries today, according to the United Nations Conference on Trade and Development. The momentum is real.

The bad news is that having a law and enforcing it are different things. And the data economy moves faster than any legislature.

What you can do, right now, is treat your own data as if you have strong protections regardless of where you live. Because the tools exist to give you that level of control, and the chapters ahead will show you how to use them. The law may not be on your side. The technology can be.

Key Privacy Frameworks Worth Knowing

1. GDPR (EU, 2018): The global benchmark. Covers all EU residents regardless of where the company processing their data is based. Includes rights of access, correction, deletion, and data portability.
2. CCPA / CPRA (California, 2020 / 2023): The strongest US state-level framework. Gives Californians the right to know, delete, and opt out of the sale of their personal data.
3. LGPD (Brazil, 2020): Closely modeled on GDPR. Applies to any company processing data of people in Brazil, regardless of where the company is based.
4. PDPA (Thailand, 2022 / Singapore, 2012, updated 2020): Both countries have established frameworks with meaningful rights, though enforcement and scope differ.
5. DPDP Act (India, 2023): Covers 1.4 billion people. Grants meaningful individual rights but includes broad government exemptions that privacy advocates continue to challenge.

Isabelle in Lyon can ask a company what it knows about her and expect a legal answer. Most people on earth cannot. That is not a technical problem. It is a political one.

CHAPTER 4

REAL PEOPLE, REAL HARM

Marco had worked for the same logistics company in Rome for eleven years. He was not wealthy by any measure, but he was careful. He paid his bills on time, kept a modest savings account, and had never missed a mortgage payment. His credit score was, as he occasionally mentioned to his wife with quiet pride, excellent.

In May 2021, he applied to refinance his mortgage and was rejected. The bank told him his credit history showed three outstanding loans he had never taken out, a mobile phone contract in his name registered to an address in Bucharest he had never visited, and a series of missed payments on a credit card he did not own.

It took Marco thirteen months, hundreds of hours, and approximately four thousand euros in legal fees to untangle. He never found out exactly how it happened. The most likely explanation, his lawyer told him, was that his personal data had been part of one of the large-scale breaches that had leaked from a data broker's database sometime in the previous two years. His name, date of birth, national identity number, and home address had been available on a dark web marketplace for less than the price of a restaurant meal. Marco is not an unusual case. He is a very ordinary one.

The Scale of the Problem

Identity theft and financial fraud enabled by stolen personal data are among the most common and most damaging consequences of the surveillance economy. The numbers are large enough to feel abstract, so let me try to make them concrete.

According to the Identity Theft Resource Center, data breaches in the United States alone exposed over 353 million individuals in 2023. That figure includes people whose records were compromised in corporate breaches, healthcare leaks, government data incidents, and financial institution failures. Each of those records typically contains enough information to open a line of credit, file a fraudulent tax return, or impersonate the individual for purposes that range from inconvenient to catastrophic.

In the UK, the Office for National Statistics recorded over 3.5 million incidents of fraud in the year ending March 2023, with the majority involving some degree of personal data misuse. The data that makes identity theft possible does not appear from nowhere. It is the same data that advertising networks collect to serve you relevant content, that loyalty programs collect to personalise your shopping experience, and that data brokers aggregate from dozens of sources to build their profiles. When that data is breached, stolen, or sold to the wrong party, the consequences land not on the company that collected it, but on the person it describes.

This is the part of the privacy conversation that tends to get lost when we focus on targeted advertising. The harm from data collection is not always abstract or slow or invisible. Sometimes it is a rejected mortgage application and fourteen months of your life spent proving you are who you always knew you were.

When Data Enables Danger

Financial harm is the entry point for this chapter because it is the most universally relatable. But it is far from the most serious harm that privacy failures enable. That territory belongs to cases where data is used not to steal money, but to locate, monitor, and control people.

The connection between data broker profiles and domestic abuse is well documented and deeply troubling. In the United States, the nonprofit Safety Net, which works with survivors of technology-facilitated abuse, has documented hundreds of cases in which abusers used commercially available data broker services to locate people who had deliberately moved to escape them. These services, designed and marketed for purposes like background checks and people searches, do not distinguish between a concerned parent looking for an estranged adult child and a violent ex-partner looking for someone who fled.

Most data brokers offer an opt-out process. The process is typically complicated, requires submission of personal identification, must be repeated separately for each broker, and needs to be redone periodically because the data tends to reappear. For a survivor of abuse who may not have a fixed address, consistent internet access, or the emotional bandwidth to navigate a bureaucratic removal process across dozens of websites, this is not a realistic solution. The burden of protection falls entirely on the person most at risk.

Location data presents a related category of danger. In 2018, a report by the New York Times demonstrated that location data purchased from a data broker could be used to track the movements of specific individuals with disturbing precision, including identifying their home, workplace, place of worship, and daily routine. The data had been collected by apps that users had downloaded for entirely unrelated purposes, apps that tracked weather, offered coupons, or provided local news. The location collection was disclosed in privacy policies. Nobody read them.

The Invisible Decisions Being Made About You

There is a category of harm that is harder to see than identity theft or stalking, but which affects far more people. It is the harm of being quietly disadvantaged by decisions made about you on the basis of data you did not know was being used.

Insurance pricing is one of the most developed applications of behavioral data outside of advertising. In several countries, insurers use data derived from credit scores, purchasing history, postcode, and in some cases social media activity to set premiums. An investigation by the Consumer Federation of America found that drivers in predominantly Black and Hispanic neighborhoods in the United States were charged higher premiums than drivers in predominantly white neighborhoods with similar or worse accident records. The variable was not risk. It was data-derived proxies for race and income.

Employment screening has developed its own data problem. Background check services in the United States hold records on hundreds of millions of people, and errors in those records are common. The 2019 Broken Records study by the National Consumer Law Center found that a significant proportion of background check reports contained inaccuracies. An error in a background check can mean a job offer withdrawn, an apartment application rejected, or a volunteer position refused. The person affected often does not know a check was run, what it contained, or that the information in it was wrong

Content recommendation algorithms present perhaps the most widespread form of this quiet harm. When a platform's algorithm determines that a user is showing signals of depression and responds by serving them more content associated with depression, not because anyone decided this was a good idea but because engagement is engagement, the result is a feedback loop that can materially affect that person's mental health. A 2021 investigation by The Wall Street Journal found that Instagram's own internal research had identified this pattern in teenage girls, showing that the platform worsened body image issues for one in three teenage girls who already struggled with them. The research was internal. The recommendation engine kept running.

None of these harms require malicious intent. They require only an economic system that optimises for data collection and engagement without accounting for the cost those activities impose on the people whose data is being collected and whose attention is being engineered.

The Chilling Effect Nobody Talks About

There is one more category of harm I want to address, and it is the most difficult to quantify but perhaps the most important to understand. It is what researchers call the chilling effect: the change in behavior that occurs when people know, or suspect, they are being watched.

A 2016 study by researcher Jonathon Penney, published in the Berkeley Technology Law Journal, found that awareness of mass surveillance programs, following the Snowden revelations in 2013, led to a measurable decline in traffic to Wikipedia articles on topics considered sensitive: terrorism, extremism, weapons, and related subjects. Traffic to those articles dropped by around 30 percent and continued falling, suggesting the effect was not temporary. People did not stop being curious. They stopped expressing that curiosity in ways that could be observed and recorded.

This matters more than it might initially appear. A functioning democracy depends on people being able to seek out information, hold unpopular opinions, explore unconventional ideas, and communicate privately. When surveillance creates a cost for doing any of these things, even a small and informal cost, the space for free thought and free expression contracts. Not dramatically, not overnight, but steadily and measurably.

The journalist Glenn Greenwald, who worked with the Snowden documents, made this point in a 2014 TED talk that has since been viewed over ten million times. His argument was simple: the claim that privacy only matters to people with something to hide fundamentally misunderstands what

privacy is for. Privacy is not about concealing wrongdoing. It is about maintaining the conditions under which people can think, speak, and act freely. Remove those conditions, and you change not just behavior but the kind of society that behavior produces.

I think about this often when I am building tools at Cyber Konekt. The goal is not to help people hide. The goal is to help people be free. These are different things, and the distinction matters.

What You Can Do About It

This chapter has covered difficult ground, and I want to end it in the same spirit as the rest of this book: with the reminder that understanding the harm is the beginning of addressing it, not a reason for despair.

The practical steps available to you fall into three broad categories. The first is reducing the data that exists about you in the first place, which is what the tools in Part Three of this book are primarily designed to help with. Less data collected means less data available to be breached, sold, or misused.

The second is monitoring what already exists. You can request your credit report in most countries for free, and checking it annually for accounts you did not open is one of the simplest and most effective early warning systems available. In the EU, you can submit subject access requests to companies and data brokers asking what they hold on you. In the US, several states give you similar rights. Using them is not paranoid. It is prudent.

The third is supporting the structural changes that individual action cannot achieve alone. Better regulation, stronger enforcement, and companies that compete on privacy rather than against it are not inevitable. They are the result of informed citizens making demands. The fact that you are reading this book is, in a small but real way, part of that process.

Marco eventually got his mortgage. His credit record was corrected. The loans that were never his were removed. The process was exhausting and expensive, and it should not have been necessary. He told his lawyer, near the end of it, that he felt like he had spent thirteen months proving he existed.

He should never have had to.

Five Steps to Reduce Your Exposure Now

1. Check your credit report. In most countries you can do this for free once a year. Look for accounts, addresses, or inquiries you do not recognise. Early detection makes resolution far less painful.
2. Search your name on two or three major data broker sites. See what they hold. Most offer an opt-out process. It is cumbersome, but worth doing. Services like DeleteMe or Incogni can automate much of this process on your behalf.
3. Review the permissions on every app that has access to your location. Revoke anything set to 'always on' that does not have a clear and immediate need for it. This reduces the location data flowing into broker pipelines daily.
4. Use a password manager and enable two-factor authentication on any account that holds financial, health, or identity information. Stolen passwords remain the most common entry point for identity theft.
5. If you are in the EU, consider submitting a subject access request to one data broker you have heard of. The response will show you, concretely, what kind of profile exists on you. It is a useful and sometimes alarming education.

The data that describes you was never meant to protect you. That has always been someone else's job. It is time to make it yours.

PART 2

HOW IT GOT THIS BAD

CHAPTER 5

**CONVENIENCE WAS
ALWAYS
THE TROJAN HORSE**

In 1994, a man named Jeff Bezos drove a car across the United States from New York to Seattle while his wife McKenzie took notes in the passenger seat. He was working out a business plan. The idea was to sell books online. Not because books were especially interesting to him, but because they were a perfect first product: uniform, easy to ship, available in vast quantities, and purchased by exactly the kind of early internet user who might trust an unknown website with a credit card number.

The genius of the plan, in retrospect, was not the books. It was the address book. Every order placed on Amazon came with a name, a delivery address, a payment method, and a record of what that person had chosen to buy. Multiply that by millions of customers and you have something worth considerably more than the margin on a paperback.

Bezos understood this early. Most of his customers did not understand it at all. They were just glad they did not have to drive to a bookshop.

This is the story of how convenience became the mechanism by which an entire generation traded away their privacy, one frictionless transaction at a time. And the most important thing to understand about it is that nobody forced anyone to do anything. The trade was always voluntary. It was also, in almost every case, invisible.

The Friction Was the Point

Before the internet, privacy was largely protected by inconvenience. Gathering information about a private individual required physical effort: visiting public records offices, making phone calls, hiring someone to follow them around. The friction was not a deliberate privacy protection. It was simply the natural cost of collecting information in an analogue world. But the effect was the same. Most information about most people was not worth the effort required to obtain it.

The internet did not just reduce that friction. It essentially eliminated it. Information that once required days of effort to assemble could now be gathered in seconds. Behavior that once left no trace began leaving an indelible one. And the companies building the infrastructure of the early web understood, faster than regulators or the public, that this shift had created something with extraordinary commercial value.

The playbook that emerged was elegant in its simplicity. Build something genuinely useful. Make it free. Collect everything. The usefulness drove adoption. The free price removed resistance. And the data collection happened quietly in the background, disclosed in legal documents that nobody read and that were not, in any meaningful sense, designed to be read.

It is worth pausing on how rational this all looked at every individual step. A consumer gets a free email account. Convenient. A company gets to show that consumer relevant advertisements. Reasonable. A third-party data broker gets

access to signals about that consumer's behavior. Legal. None of these steps, taken in isolation, sounds alarming. The alarm is only audible when you zoom out and look at what the system as a whole produces: a surveillance infrastructure of extraordinary reach, built with the enthusiastic participation of the people being surveilled.

How They Made You Want It

Convenience alone does not explain the scale of what happened. Plenty of convenient things fail to achieve mass adoption. What the technology companies of the 2000s and 2010s understood, and invested heavily in understanding better, was behavioral psychology.

The slot machine is the standard reference here, and it has become a cliché, but it became a cliché because it is accurate. Variable reward schedules, the mechanism by which a slot machine pays out unpredictably rather than consistently, are among the most powerful tools for sustaining repetitive behavior that psychology has identified. The social media feed, with its unpredictable mixture of content that might be interesting or boring, validating or disappointing, is a variable reward schedule running at scale on a device that most people carry everywhere they go and check, on average, dozens of times per day.

Aza Raskin, the designer widely credited with popularising the infinite scroll, the mechanism by which social media feeds load new content automatically as you reach the bottom, has spoken publicly about his regret at creating it. He has estimated that infinite scroll wastes around 200,000 human

lifetimes every day. He invented it in about an hour. He did not charge extra.

The notification was another piece of the architecture. Not the useful notification that tells you your flight is delayed or your package has arrived, but the social notification: someone liked your post, someone replied to your comment, someone tagged you in a photo. Each of these is a small trigger, engineered to pull you back to the platform at the moment you might otherwise have put your phone down. The data generated by your return is, of course, the point.

None of this was hidden, exactly. It was discussed openly in design circles, written about in books like *Hooked* by Nir Eyal, and taught in product design courses at universities. The techniques for engineering habitual product use were not secret. They were a discipline. What was less openly discussed was what happened to the behavioral data those habits generated, and who profited from it.

The Consent That Was Not Really Consent

Every few years, a researcher runs an experiment in which they offer people small rewards, usually a pizza or a few dollars, in exchange for signing a document that includes, buried in its terms, a clause granting the researcher their immortal soul, the rights to their firstborn child, or some other clearly absurd concession. A reliable majority of participants sign without noticing.

This is funny until you apply it to the terms of service documents that govern your relationship with every significant digital platform you use. These documents are not designed to be read. The average terms of service agreement for a major platform runs to somewhere between ten and thirty thousand words, written in legal language at a reading level calibrated for professionals, and presented at a moment when the user's goal is to get past the document and use the product as quickly as possible.

A study by Carnegie Mellon researchers Lorrie Faith Cranor and Aleecia McDonald, now somewhat dated but never seriously challenged in its conclusions, estimated that if Americans actually read every privacy policy they encountered in a year, it would consume approximately 76 work days of reading time per person. The entire framework of informed consent on which the surveillance economy's legal legitimacy rests is built on the assumption that people can and do make meaningful choices about their data. They cannot, and they do not, and the companies writing those policies know this better than anyone.

The GDPR attempted to address this with its requirement that consent be freely given, specific, informed, and unambiguous. It was a genuine improvement. It is also why you now spend three seconds every morning clicking through cookie banners on every website you visit, a process that has become so universal and so universally resented that it has arguably made people less thoughtful about privacy rather than more. When every website asks for your consent to everything all the time, consent stops feeling meaningful.

That may not have been accidental.

The Gradual Ratchet

One of the most effective mechanisms by which the surveillance economy expanded was gradualism. The amount of data collected from users, and the ways it was used, did not leap from zero to everything overnight. It moved in small steps, each of which seemed reasonable in the context of the previous step, none of which would have been acceptable if proposed at the beginning.

Consider the trajectory of a single major platform over two decades. It begins by asking for your name and email address to create an account. Then it asks for your date of birth, ostensibly for age verification. Then it begins tracking which posts you engage with, to show you more relevant content. Then it introduces a mobile app that, with your permission, accesses your location. Then that location access becomes more persistent. Then it shares behavioral signals with advertising partners. Then those partners share data with data brokers. Then the data broker sells your profile to an insurer.

At no single point in that chain did anything dramatically cross a line that had not already been crossed. Each step was normalised by the previous one. The ratchet moved in one direction only, and it moved slowly enough that most people never felt it turn.

I have watched a version of this dynamic play out in the privacy tools industry too. The companies that genuinely protect your privacy compete in a market where the free

alternatives are funded by the data they collect from you. Every time a user chooses the free option, they make the privacy-protecting alternative slightly less viable and the surveillance-funded one slightly more dominant. Convenience, again, does the work that no amount of hostile persuasion could.

What Was Actually Being Traded

I want to be fair here, because this chapter risks sounding like a straightforward condemnation of an industry that has, in genuine ways, made life better for a lot of people.

Free email changed how the world communicates. Free maps ended the era of driving in circles for forty-five minutes before admitting you are lost. Free search made the collective knowledge of humanity available to anyone with a phone. Free translation tools crossed language barriers that once required years of study. These are not trivial benefits. They are real, and they were delivered, mostly, to people who could not have afforded to pay for them.

The question is not whether the trade produced anything of value. It did. The question is whether the people making the trade understood what they were trading, and whether they would have agreed to those terms if they had. I think the honest answer, in most cases, is no. Not because people are foolish, but because the terms were never presented clearly, the consequences were not foreseeable at the point of exchange, and the power asymmetry between a person clicking 'agree' and a company with a thousand engineers optimising for data collection was never going to produce a genuinely

balanced negotiation.

The good news is that awareness changes the calculation. Once you understand the terms of the trade, you can start making different choices. Not all at once, and not without cost, but meaningfully. The chapters that follow are largely about how.

The One Thing You Can Do Right Now

Open the settings of your most-used social media app and find the section for data or privacy. Look specifically for the option to download your data, which most major platforms are now legally required to offer. Request it.

What arrives, usually within a day or two, will be a compressed file containing everything the platform holds about you. Your posts and messages, yes. But also your search history within the app, every advertisement you were shown and whether you clicked it, the precise times you opened and closed the app, your inferred interests and demographic categories, and in many cases data points you have no memory of providing.

It is one of the more clarifying experiences available in the modern world. Not comfortable, exactly. But clarifying. And understanding what has already been collected is the first step toward deciding what you want collected going forward.

Five Signs a Free Product Is Expensive in Ways You Have Not Noticed

1. The product asks for permissions it does not obviously need. A torch app that wants your contacts list is not offering you a torch. It is offering you a contacts list a torch.
2. The privacy policy is longer than this chapter. This is not a coincidence. Length is a deterrent, not a disclosure.
3. The product gets better the more you use it, in ways that require it to learn about you. Useful, certainly. But ask what it is learning, and where that learning goes.
4. The company's revenue model is not immediately obvious. If you cannot quickly identify what customers pay for, you are probably looking at a product where the data is the revenue model.
5. Opting out of data collection makes the product significantly less functional. This is the technical implementation of 'agree or leave,' dressed up as a feature.

The Trojan horse worked because the people of Troy wanted what was inside it. That has always been the point.

CHAPTER 6

**THE FINGERPRINTING
ARMS RACE**

Imagine you walk into a shop, buy nothing, and leave. The next day you walk into a completely different shop across town. The owner greets you by name, knows roughly how much money you earn, remembers that you spent eleven minutes browsing winter coats yesterday, and has already selected three options she thinks you will like based on that.

You never gave either shop your name. You paid cash. You did not sign up for anything. You simply walked in and walked out. This is roughly what happens every time you visit a new website. And the mechanism that makes it possible is something most people have never heard of: device fingerprinting.

Understanding fingerprinting does not require any technical knowledge. It requires only the recognition that your device, the phone or laptop you are reading this on, is not a blank, anonymous object. It is a remarkably distinctive one. And that distinctiveness is being read, right now, by systems designed to turn it into a permanent, trackable identity.

What a Fingerprint Actually Is

Every device you own has a personality. Not in the sentimental sense, but in the technical one. The screen you are reading this on has a specific resolution. Your browser is a specific version of a specific application. Your device's clock is set to a specific timezone. The list of languages your browser accepts, the fonts installed on your system, the way your graphics processor renders a particular image: all of these details are slightly different from device to device.

None of these details, taken individually, is especially identifying. Millions of people use Chrome. Millions of people are in the same timezone. But combine enough of them together and something remarkable happens. The combination becomes unique. Research using the AmIUnique platform, which analysed 118,934 browser fingerprints across 17 measured attributes, found that 81 percent of mobile fingerprints were unique. No name required. No login. No cookie. Just the characteristics of the device itself, assembled into a profile as distinctive as a thumbprint.

This is a fingerprint. And unlike a cookie, which you can delete, a fingerprint cannot be cleared. It is not stored on your device. It is calculated fresh each time you visit a site, from information your device broadcasts automatically simply by connecting to the internet. Deleting your browsing history does nothing to it. Opening a private window does nothing to it. The fingerprint persists because the device does.

Fingerprinting is currently found on more than a quarter of the top ten thousand websites in the world, according to research published in the journal of the Association for Computing Machinery. Most users on those sites have no idea it is happening. There is no notification. No cookie banner. Nothing to click through or decline.

The Cookie Was Never Really the Problem

For most of the past decade, the public conversation about online tracking has focused on cookies. Cookies are small files that websites store on your device to remember who you are. They are useful for keeping you logged into your email account

and remembering what you put in your shopping basket. They are also widely used to track you across websites, building a picture of your browsing behavior that advertisers pay to access.

Cookies became a public issue because they are visible and deletable. Regulators could point to them. Laws could be written about them. Users could, with some effort, clear them. The GDPR in Europe and similar laws in other regions required websites to get your consent before placing tracking cookies. The result was the cookie banner: that slightly aggressive popup on every website you visit asking you to accept, decline, or customise settings that most people have never read.

The advertising industry watched all of this with something between concern and patience. Concern, because cookies were genuinely useful. Patience, because they already knew what would replace them.

Google announced in 2020 that it would phase out third-party tracking cookies from its Chrome browser entirely. This was presented as a significant privacy win. Chrome is used by roughly two thirds of the world's internet users, so a cookie phase-out would have genuinely reduced a major tracking mechanism. The advertising industry, which had spent years building its infrastructure on cookies, immediately began investing in alternatives. Fingerprinting was at the top of that list.

In July 2024, Google reversed course entirely. Rather than eliminating cookies, it announced it would introduce a user-choice model, letting Chrome users decide their own tracking preferences. The advertising industry, which had been loudly opposing the cookie phase-out since 2020, breathed a collective sigh of relief. Privacy advocates were less enthusiastic. But the episode illustrated something important: even when the infrastructure of tracking comes under genuine pressure, the industry's incentive to find workarounds is enormous, well-funded, and fast.

Why Incognito Mode Will Not Save You

This is the section most people find uncomfortable, so I want to be straightforward about it.

Private browsing modes, called incognito in Chrome, private in Firefox, and various names in other browsers, do one thing: they prevent your device from storing a local record of your browsing session. Your history is not saved. Cookies placed during the session are deleted when you close it. If someone else picks up your device and opens the browser, they will not see what you were doing.

What private mode does not do is make you invisible to the websites you visit. Your fingerprint is unchanged. Your IP address is unchanged. The fingerprinting scripts running on the sites you visit function identically in private mode. The advertising networks tracking your behavior across sites continue tracking it. Private mode protects you from your device. It does not protect you from the internet.

I say this not to be discouraging, but because understanding

the actual limits of a tool is the only way to use it appropriately. Private mode has genuine uses. It is useful for keeping a surprise purchase out of your shared device's history. It is useful if you are using a computer that is not yours. It is not, and was never designed to be, an invisibility cloak.

The gap between what people believe private mode does and what it actually does is one of the largest misconceptions in everyday digital life. A 2017 study by researchers at Chicago University and Leibniz University Hannover found that a significant majority of users believed private mode hid their activity from their internet provider, their employer, and the websites they visited. None of those things are true.

How the Race Gets Run

I have spent years building tools designed to resist fingerprinting, and I want to give you an honest account of what that work has taught me.

The arms race between tracking and privacy protection is real, and it moves fast. When browsers began blocking certain fingerprinting methods, trackers moved to new ones. When canvas fingerprinting, a technique that identifies your device by how it renders a tiny invisible image, became well known and began to be blocked, the industry developed audio fingerprinting, which does the same thing using your device's sound processing. When that became detectable, new techniques emerged using the precise behavior of your graphics hardware.

Each round of the race follows the same pattern. Privacy tools and browsers identify and block a tracking method. The tracking industry develops a new one that works around the block. Privacy tools catch up. The cycle repeats. The gap between a new fingerprinting technique being deployed and privacy tools being able to detect and counter it is measured in weeks to months, during which millions of users are being tracked by a mechanism their protection tools do not yet know about.

The honest conclusion from years of working in this space is that no single tool makes you invisible. What good privacy tools do is raise the cost and complexity of tracking you, push you out of the easy, automated tracking systems that capture the vast majority of behavioral data, and ensure that the data collected about you is significantly less complete and less accurate than it would otherwise be. That matters. Imperfect privacy is not useless privacy. But it is worth being clear about what the tools can and cannot guarantee.

This is also why the combination of tools matters more than any individual one. A VPN changes your apparent location and hides your traffic from your internet provider. A privacy-focused browser with fingerprint protection makes your device harder to uniquely identify. A private search engine ensures your queries are not being logged against your profile. None of these is a complete solution on its own. Together, they reduce your trackable surface area substantially.

The One Thing That Surprises Almost Everyone

There is a dimension of fingerprinting that even people who consider themselves reasonably privacy-aware tend not to know about. It is called a behavioral fingerprint, and it has nothing to do with your device's technical characteristics.

Research published in the journal *Nature* found that knowing just four websites a person visits regularly is enough to identify them with 95 percent accuracy from a large anonymised dataset. Not their name. Not their device. Just four websites. Your browsing pattern, the particular combination of sites you visit in a typical week, is as distinctive as your hardware configuration. Perhaps more so, because it reflects your actual habits, interests, and routine in a way that technical device characteristics do not.

This means that even if you successfully disguised your device's fingerprint entirely, the pattern of your behavior online would still tend to identify you. Tracking does not require any single piece of individually identifying information. It requires only enough data points that, in combination, only one person plausibly fits the description.

Understanding this changes how you think about privacy tools. The goal is not to become invisible. The goal is to become less distinctive: to reduce the number and precision of the data points available, to introduce enough variation and noise into your digital behavior that the profile assembled about you is incomplete, inaccurate, or economically not worth the effort of maintaining.

That is an achievable goal. And Part Three of this book is largely about how to work toward it.

Five Things You Can Do to Reduce Your Fingerprint

1. Use a browser with built-in fingerprint protection. Brave and Firefox with privacy settings enabled both actively work to make your fingerprint less unique by returning standardised values to fingerprinting scripts rather than your real device characteristics.
2. Install uBlock Origin. It blocks many of the scripts used for fingerprinting before they can run. It is free, effective, and used by tens of millions of people. One of the most reliable single steps available.
3. Use a VPN consistently. It does not stop fingerprinting, but it changes your apparent location and prevents your internet provider from seeing your browsing behavior, which removes one layer of the tracking stack.
4. Avoid logging into accounts on sites where you want to browse anonymously. An account login collapses all the uncertainty of fingerprinting into a confirmed identity immediately.
5. Vary your browsing patterns occasionally. This is a small step but a genuine one. Behavioral fingerprinting depends on predictability. Unpredictability, even modest amounts of it, degrades the accuracy of behavioral profiles.

The shop owner who greeted you by name did not recognise your face. She recognised the way you walk.

CHAPTER 7

WHY THE RULES ARE NOT WORKING

In 2024, Ireland's Data Protection Commission, the body responsible for regulating most of the world's largest technology companies under the GDPR, issued fines totalling 652 million euros. LinkedIn was fined 310 million euros. Several other major platforms received substantial penalties. On paper, it looks like serious enforcement.

Of that 652 million euros, 582,500 was actually collected. Not 582 million. Not 58 million. Five hundred and eighty-two thousand and five hundred euros, out of 652 million levied. Less than one tenth of one percent. The rest is being appealed, contested in court, or simply sitting uncollected while legal proceedings continue across multiple jurisdictions. Since the GDPR came into force in 2018, more than four billion euros in fines have been issued by Irish regulators alone. The amount actually paid is a small fraction of that.

This is the gap that matters. Not the rules as written, but the rules as enforced. And it is the gap that explains why, seven years after the world's most ambitious privacy regulation came into force, the surveillance economy is larger than it has ever been.

The Problem With Being the World's Privacy Police

The GDPR's one-stop-shop mechanism was designed to simplify compliance for international companies. Rather than navigating separate regulatory processes in each of the EU's member states, a company with a European headquarters in one country would be supervised primarily by that country's regulator. Efficient in theory. Awkward in practice.

Most of the world's largest technology companies chose to base their European operations in Ireland. Low corporate tax rates, an English-speaking workforce, and a generally business-friendly environment made it an attractive choice. The side effect, unintended but substantial, was that Ireland became the primary data protection regulator for companies like Meta, Google, Apple, Twitter, LinkedIn, and TikTok.

Ireland is a country of five million people. It is being asked to regulate companies whose combined market capitalisation exceeds the GDP of most continents. Its Data Protection Commission has faced sustained criticism from privacy advocates, other EU member state regulators, and the European Data Protection Board for moving too slowly, resolving too few complaints, and, until recently, imposing too few significant penalties.

The DPC disputes this characterisation. The gap between billions levied and thousands collected does not resolve the dispute in its favour.

This is not a criticism of Ireland specifically. It is a structural problem baked into the regulation's design. When the entity responsible for policing the world's most powerful companies is a small national regulator whose budget the Irish parliament has repeatedly failed to fully fund, the outcome should not be surprising. Good rules, enforced inadequately, produce the appearance of accountability without much of the substance.

The United States: Where Good Bills Go to Wait

In July 2022, the American Data Privacy and Protection Act passed the House Energy and Commerce Committee with a vote of 53 to 2. Bipartisan. Near-unanimous. The closest the United States had ever come to passing a comprehensive federal privacy law.

It never made it to a floor vote in either chamber. By January 2023, when the congressional session ended, the bill had expired without being formally considered. A successor bill, the American Privacy Rights Act, was introduced in 2024. It also stalled.

The reasons are instructive. The ADPPA fell apart over two specific disagreements. The first was state preemption: whether a federal law should override stronger state-level protections, particularly California's, or whether states should be free to set higher standards. California's representatives were unwilling to see their stronger protections diluted by a federal floor. The second was the private right of action: whether individual citizens should be able to sue companies directly for privacy violations, or whether enforcement should be left to regulators and state attorneys general. Democrats largely supported individual lawsuits. Republicans largely opposed them.

Neither of these disagreements is trivial. They reflect genuine and reasonable differences in legal philosophy.

But the practical result is that the United States, home to the world's most influential technology companies and the world's largest digital advertising market, remains without a single federal privacy law. Individual states fill the gap unevenly. Companies navigate the patchwork with varying degrees of compliance and creativity. And the absence of a floor means that the majority of Americans have no guaranteed federal right to know what data is held about them, correct it, or have it deleted.

Self-Regulation: The Fox and the Henhouse

In the absence of strong external regulation, the technology industry has periodically offered to regulate itself. Privacy frameworks, voluntary commitments, transparency reports, and industry codes of practice have accumulated over the past two decades into an impressive-looking body of self-governance that has done relatively little to change the underlying dynamics of the surveillance economy.

This is not because the people writing these commitments are dishonest. Most are not. It is because self-regulation has a structural flaw that no amount of good intention can resolve: the entity doing the regulating and the entity being regulated have identical interests. A company that profits from data collection will, when given the choice, interpret ambiguous situations in ways that favour data collection. It will invest in compliance theatre, the visible apparatus of responsible data stewardship, while preserving the practices that generate revenue. This is not cynicism. It is a predictable response to incentives.

The history of the advertising technology industry is full of examples. Companies that publicly committed to not tracking users across websites built technical systems that achieved exactly the same result through different mechanisms.

Companies that promised not to sell user data sold access to audiences defined by that data, a distinction with no practical meaning for the user. Companies that offered opt-out mechanisms made them difficult to find, complicated to use, and ineffective when used, because the opt-out applied only to the company itself and not to the dozens of partners it had already shared data with.

The counterargument, which is worth taking seriously, is that some self-regulation is better than none, that voluntary standards raise the floor even when they fail to reach the ceiling, and that the reputational cost of being caught violating a public commitment creates genuine incentives for compliance. This is true. But it is also true that the gap between the voluntary commitments and the actual practices of the industry over the past twenty years has been large, consistent, and well-documented.

The Speed Problem

There is a structural mismatch at the heart of privacy regulation that no legislative body has yet solved: technology moves faster than law.

The GDPR was drafted between 2012 and 2016, came into force in 2018, and was based on a model of data processing that was already becoming outdated when the ink dried. It was not

written for artificial intelligence systems that infer sensitive attributes from apparently innocuous data. It was not written for large language models trained on scraped personal information. It was not written for the real-time bidding systems that share your personal data with hundreds of companies in the milliseconds before a webpage loads, a process that the GDPR's consent requirements were arguably never equipped to govern.

Regulators are now scrambling to apply seven-year-old rules to technologies that did not exist in their current form when those rules were written. The European Data Protection Board has issued guidance on AI. Individual member state regulators have opened investigations into generative AI systems. The pace is, by the standards of technology development, glacial.

This is not a criticism of the regulators, most of whom are working with genuine commitment and limited resources. It is a description of a fundamental asymmetry. The companies being regulated employ thousands of engineers and lawyers. The regulators employ hundreds of people, many of whom are not technical specialists, and operate on government budgets that have not kept pace with the scale of what they are being asked to oversee. The contest is not fair, and the outcomes reflect that.

What Good Regulation Would Actually Require

I want to be clear that I am not arguing regulation cannot work. The GDPR, for all its enforcement problems, has changed corporate behavior in measurable ways. The cookie banner, for all that it has become a frustrating ritual, represents a genuine shift in the legal baseline. The right to access your data and request its deletion, however imperfectly implemented, did not exist for most people a decade ago.

What I am arguing is that good privacy regulation requires several things that most current frameworks lack. It requires adequate funding. A regulator that cannot afford to investigate a company is not a regulator. It is a suggestion. The Irish DPC's budget has been chronically insufficient relative to its caseload. Most national data protection authorities face similar constraints.

It requires technical expertise. The gap between what regulators understand about the systems they are regulating and what the companies understand about those systems is wide and growing. Closing it requires investment in people with deep technical knowledge, which means paying salaries competitive with the private sector. Most public bodies cannot do this.

It requires meaningful penalties that are actually collected. A fine that sits uncollected while appeals proceed for five years is not a deterrent. It is a line item on a balance sheet. Effective enforcement requires mechanisms to ensure penalties are paid, not merely levied.

And it requires international coordination. Data crosses borders in milliseconds. Regulation crosses them slowly, through treaties, mutual recognition agreements, and diplomatic negotiation. The mismatch between the speed of data flows and the speed of regulatory cooperation is one of the largest unresolved problems in global privacy governance.

None of these requirements is impossible to meet. Several countries are making genuine progress. But the honest assessment is that as of today, the rules are not working well enough, fast enough, or consistently enough to keep pace with the surveillance economy they are trying to govern.

What You Can Do While Waiting for the Rules to Catch Up

This chapter has been heavier going than most, and deliberately so. Understanding why the rules are failing matters, because it explains why individual action remains important even as regulatory frameworks improve. If the law fully protected you, this book would be shorter.

The practical implication is this: do not wait for regulation to solve the problem on your behalf. Use the tools available to you now. Support organisations that are pushing for better enforcement, stronger legislation, and adequately funded regulators. And when you encounter a company that makes your privacy a genuine priority rather than a compliance exercise, choose it over one that does not.

The rules will improve. Enforcement will become more consistent. International cooperation will slowly become more effective. But the gap between now and that future is measured in years, possibly decades, and the data being collected about you in the meantime will still exist long after the rules that should have prevented its collection are finally in place.

Five Ways to Tell If a Company Takes Privacy Seriously

1. It publishes a transparency report. Companies that genuinely respect privacy tend to report publicly on government data requests, the number of user complaints they receive, and how those complaints are resolved.
2. Its privacy settings are easy to find and easy to use. Genuine privacy commitments do not hide behind seven nested menus. If finding the opt-out requires a search engine, treat that as information.
3. It collects only what it needs. A notes app does not need your location. A weather app does not need your contacts. Companies that ask for permissions beyond their obvious function are telling you something about their business model.
4. It responds to data access requests promptly. Under GDPR and equivalent laws, you can ask companies what data they hold on you. Companies with good data practices can answer this quickly. Those without good practices often cannot.
5. It does not rely on advertising revenue. The incentive structure of advertising-funded services is fundamentally in tension with genuine privacy protection. This does not make every ad-funded service untrustworthy, but it is a constraint worth keeping in mind.

The fine was 652 million euros. The amount collected was 582,000. The law exists. Whether it functions is a different question entirely.

PART 3

WHAT YOU CAN DO

CHAPTER 8

THE LAYERED DEFENSE

Author's note: Cyber Konekt has affiliate relationships with some tools recommended in this chapter, including NordVPN and Surfshark. These relationships are fully disclosed in the author's note at the back of this book. They did not influence which tools are recommended here

There is a particular kind of paralysis that sets in when people start learning about online privacy. It usually happens somewhere around the third article about fingerprinting, or the second paragraph about cookies, or the first time someone explains that deleting your browsing history is roughly as effective as closing your curtains to hide from a satellite. The problem feels so large, and the solutions feel so technical, that doing nothing begins to seem like a reasonable response.

It is not a reasonable response. But I understand why it happens, and I want to spend this chapter dismantling it.

The truth is that you do not need to understand how tracking systems work in order to reduce their reach. You do not need to be a developer, or a security researcher, or someone who reads technical documentation for fun. What you need is a small set of tools, layered sensibly, and a basic understanding of what each one does and does not do. That is genuinely achievable in a single afternoon, and the effect it has on the amount of data collected about you is substantial.

Think of it like locking your front door. You probably do not know how the lock mechanism works. You do not need to. You just need to know that locking it makes your home meaningfully harder to enter than leaving it open.

Privacy tools work the same way. Imperfect, yes. But far better than nothing.

Why Layers Matter More Than Any Single Tool

The surveillance ecosystem described in earlier chapters is not a single system. It is a stack of overlapping systems, each collecting data through different mechanisms. Your internet provider sees which websites you visit. The websites themselves see your fingerprint and your behavior. The advertising networks embedded in those sites see your cross-site history. The apps on your phone see your location, your contacts, and your usage patterns. Your device manufacturer may see all of it.

No single privacy tool addresses all of these layers simultaneously. A VPN hides your traffic from your internet provider but does nothing about fingerprinting. A privacy-focused browser reduces fingerprinting but cannot prevent your phone's apps from collecting location data. An ad blocker stops many tracking scripts but does not encrypt your connection. Each tool addresses a specific part of the problem. The power comes from using them together.

This is what I mean by a layered defense. Not a single wall, but a series of overlapping protections, each one covering the gaps left by the others. The goal is not to achieve perfect invisibility, which is neither possible nor necessary. The goal is to make the data profile assembled about you substantially less complete, less accurate, and less valuable than it would be

with no protection at all. At that point, you shift from being an easy target to being an expensive one. Most tracking systems will simply move on.

Layer One: The VPN

Approximately 1.75 billion people worldwide used a VPN in 2025. That number is growing fast, driven partly by privacy awareness and partly by the practical need to access content across geographic restrictions. A VPN has become, for many people, the first privacy tool they try.

What a VPN actually does is straightforward. It creates an encrypted tunnel between your device and a server operated by the VPN provider. Your internet traffic passes through that tunnel before going anywhere else. From the perspective of your internet provider, all they see is that you are connected to a VPN server. They cannot see which websites you are visiting, what you are downloading, or what you are doing online. From the perspective of the websites you visit, your apparent location is wherever the VPN server is, not where you actually are.

What a VPN does not do is equally important to understand. It does not make you anonymous. The VPN provider itself can see your traffic, which is why choosing a trustworthy one matters enormously. It does not stop fingerprinting. It does not prevent websites from tracking your behavior once you are on them. And it does not protect you from the apps on your phone, which typically bypass the VPN entirely for certain types of data.

The free VPN problem is worth addressing directly. A 2024 study conducted by Top10VPN's Head of Research Simon Migliano, found that 88 percent of free Android VPNs leak user data. Many free VPN providers fund their operations by logging your browsing behavior and selling it to data brokers, which is precisely the outcome you were trying to avoid. If a free VPN has no clear business model, using it for privacy is roughly analogous to asking a data broker to hold your private documents for safekeeping. The ones worth using are the ones with independently audited no-logging policies and a clear, subscription-based business model. The cost is typically a few euros per month. It is worth it.

When choosing a VPN, look for one that has had its no-logging policy independently audited by a reputable third party, not just one that claims to have such a policy. Several established providers have now undergone multiple independent audits, which is the standard worth looking for. Jurisdiction matters too: a VPN provider based in a country with strong privacy laws and no mandatory data retention requirements is a meaningfully different proposition from one based somewhere that requires it to cooperate with government requests.

Layer Two: The Browser

Most people use Chrome. This is understandable. Chrome is fast, well-designed, and deeply integrated with services that many people use daily. It is also developed by a company whose revenues are substantially linked to advertising. Advertising-funded services often rely on data about user behaviour. These facts are not unrelated.

The browser you use is one of the most consequential privacy decisions you make, because it is the lens through which you experience almost everything on the internet. A privacy-focused browser does several things a standard one does not. It blocks tracking scripts by default. It resists fingerprinting by returning standardised information to fingerprinting systems rather than your real device characteristics. It upgrades connections to secure versions automatically. And it does all of this without requiring you to install anything extra or configure anything manually.

Brave has grown its user base by over 21 percent year on year and has become the most feature-complete privacy browser for most everyday users. It blocks ads and trackers aggressively by default, resists fingerprinting, and integrates Tor for situations where you want a higher level of anonymity. Firefox remains the strongest choice if you want extensive customisation and extension support, though it requires more deliberate configuration to reach its privacy potential out of the box. DuckDuckGo's browser, which has surpassed 50 million all-time downloads, offers the simplest experience for users who want strong default privacy without any configuration at all.

Switching browsers takes about ten minutes and costs nothing. Of all the steps in this chapter, it is probably the one with the highest ratio of privacy improvement to effort required.

Layer Three: The Search Engine

Every search you make is a window into what you are thinking about, worried about, curious about, and planning to do. Major search engines log these queries, associate them with your

profile, and use them to build a behavioral model that informs the advertisements you see and the content you are served. A single search for symptoms of a medical condition can trigger health-related advertising that follows you across the web for months.

Privacy-focused search engines work differently. They do not log your queries. They do not associate searches with a personal profile. They do not personalise results based on your history. The search results are, in most cases, good enough for everyday use, and the gap between them and the results from a major search engine has narrowed significantly over the past few years.

DuckDuckGo is the most widely used privacy search engine and a reasonable default for most people. Brave Search operates its own independent search index rather than relying on another company's results, which makes it more genuinely independent. Startpage shows Google results without passing any identifying information to Google, which is useful if you specifically value Google's search quality but not its data collection. Each has trade-offs, and switching your default takes about thirty seconds in any browser's settings.

Layer Four: The Ad Blocker

Ad blockers have an image problem. They are often framed as tools for avoiding paying for content, which misses most of what they actually do. A modern ad blocker does not just hide banner advertisements. It blocks the tracking scripts, third-party analytics tools, social media buttons, and fingerprinting code that are embedded in the majority of websites you visit.

It is, in practical terms, as much a privacy tool as an advertising one.

uBlock Origin is the most widely recommended option, and for good reason. It is free, open source, extremely effective, and imposes almost no performance cost on your browsing. It runs in the background, blocking tracking requests before they can run, without requiring any configuration. For most people it is a set-once-and-forget tool that quietly reduces the data flowing out of their browser with every page they load.

One note worth making here: if you switch to Brave, it has an ad blocker built in that is effective enough for most purposes. You do not necessarily need a separate extension. But if you stay with another browser, adding uBlock Origin is one of the most impactful single steps you can take.

The Stack, Assembled

Put these four layers together and you have a meaningfully different relationship with the internet than most people have. Your internet provider cannot see what you are doing. Tracking scripts cannot run freely in your browser. Your searches are not logged against your profile. The fingerprinting systems on the websites you visit are receiving standardised noise rather than your real device characteristics. The data profile being assembled about you is substantially less complete than it would have been an hour ago.

I want to be honest about what this does and does not achieve. It does not make you invisible. If you log into an account, that account's provider knows you are there. If you use a mobile

app, the app may collect data through channels your browser protection does not reach. If you are the target of a determined, well-resourced surveillance effort, a consumer privacy stack will not stop it. These tools are not designed for that scenario.

What they are designed for is the scenario that describes almost everyone reading this book. They reduce your exposure to the automated, commercial surveillance systems that operate at scale across the entire internet, treating every connected person as a data source to be processed. That is the threat model that matters for most people's daily lives. And for that threat model, the layered approach described in this chapter is genuinely effective.

I have built tools in this space for years, and the honest thing I can tell you is this: the gap between someone who uses a sensible privacy stack and someone who uses none is large. Not perfect versus imperfect. Large versus small. That gap is worth closing. It is precisely because assembling and managing these layers felt unnecessarily complicated for most people that we built VPN Toolkit: a single Android app that bundles the core layers of a privacy stack, covering VPNs, browsers, search engines, and ad blockers, into one place, so that the decision is made once rather than four separate times across four separate apps.

Starting Today, Not Someday

The reason most people do not do any of this is not that they do not care. It is that the moment of caring and the moment of acting are almost always separated by enough time for inertia

to win. So here is a deliberately short version of the stack, in the order that gives you the most benefit for the least effort. First, switch your browser to Brave or Firefox this afternoon. It costs nothing and takes ten minutes. Second, change your default search engine to DuckDuckGo or Brave Search in your new browser's settings. Thirty seconds. Third, if you do not already have a VPN, choose one with an independently audited no-logging policy and set it to connect automatically. If you are already using a free one, consider whether it is actually protecting you or doing the opposite. Fourth, if you are using Firefox, install uBlock Origin. If you are using Brave, the built-in blocker handles most of what you need.

That is it. Four steps, one afternoon, a substantially different privacy posture. The chapters that follow build on this foundation with more specific protections for your files, your communications, and the different contexts of your digital life.

But the stack described here is the foundation everything else rests on. Start here.

Your Privacy Stack at a Glance

VPN: Choose one with an independently audited no-logging policy. Pay for it. NordVPN, Surfshark VPN, and Mullvad are among those that have passed multiple independent audits. Avoid free VPNs unless from a well-established provider with a clear freemium model.

Browser: Brave for most users who want strong defaults with no configuration. Firefox for users who want maximum control and customisation. Either is a significant improvement over Chrome for privacy.

Search engine: DuckDuckGo as a reliable, widely used default. Brave Search for independence from any major company's index. Startpage if you specifically want Google-quality results without Google's tracking.

Ad blocker: uBlock Origin if you use Firefox or another browser. Brave's built-in shield if you use Brave. Either blocks the tracking scripts that follow you across the web.

Bonus layer: Set your DNS to a privacy-respecting provider such as Cloudflare's 1.1.1.1 or Quad9. DNS requests reveal every domain you visit to your internet provider by default. A privacy-respecting DNS resolver stops that specific leak without affecting anything else.

You do not need to understand the lock to know that turning the key makes things safer. You just need to turn the key.

CHAPTER 9

YOUR FILES, YOUR COMMUNICATIONS

Think about what is currently sitting in your phone's photo roll. Not the holiday pictures and the screenshots of recipes you never made. The other things. The photograph of your passport you took before a trip. The picture of your medical test results. The document you signed and photographed because you did not have a scanner. The screenshot of a bank statement. The images of the things that, if they ended up somewhere they should not be, would cause you genuine harm.

Most people's photo rolls are a remarkably complete archive of the most sensitive moments of their lives, stored in an app connected to a cloud service they set up in five minutes five years ago, protected by whatever default settings were in place at the time, which they have never checked since.

The average cost of a data breach reached 4.88 million US dollars in 2024, according to IBM's annual Cost of a Data Breach report, the largest single-year increase since the pandemic. That figure describes the cost to organisations. The cost to the individuals whose files were in those breached systems is harder to quantify and never makes the headline. This chapter is about your files and your communications: what happens to them, who can see them, and what genuinely protects them. The answers are more accessible than most people assume, and the gap between doing nothing and doing something is meaningful.

The File on Your Phone Is Not as Safe as It Feels

When you take a photograph on a modern smartphone, several things happen simultaneously. The image is saved to your device's local storage. If you have cloud backup enabled, which most people do because it was the default when they set up the phone, a copy is uploaded to a server owned by a technology company. That company stores it, indexes it, and in many cases runs it through automated analysis to identify faces, objects, and locations.

This is not necessarily sinister. Cloud backup serves a genuinely useful purpose: it means you do not lose years of photographs when your phone is stolen or drops into a lake. But it also means that a file you thought was sitting privately on your device is in fact sitting on someone else's server, subject to their security practices, their policies, their legal obligations, and their business interests.

When the security practices fail, as they regularly do, your files are exposed. In 2024, a breach at a major cloud data company called Snowflake compromised data held by dozens of its corporate clients, including AT&T and Ticketmaster, in what one researcher described as one of the largest breach events of the year. The vector was simple: stolen login credentials. No sophisticated attack. Just a username and password that should not have been in the attacker's hands.

The point is not that cloud storage is inherently unsafe. It is that trusting your most sensitive files to default settings, without understanding where they go or how they are protected, is a choice with consequences. Understanding those consequences takes about ten minutes. Making different choices takes a little longer.

What Encryption Actually Does to a File

Encryption is one of those words that appears constantly in privacy writing and is rarely explained in a way that means anything to a non-technical reader. Let me try.

Imagine you have a letter. You put it in a locked box. You give the box to a courier who takes it across the city. Anyone who intercepts the courier can take the box, but they cannot read the letter, because the box is locked and only you and the intended recipient have the key. Encryption is that locked box, applied to digital files.

When a file is encrypted, it is transformed into a version that looks like meaningless noise to anyone who does not hold the decryption key. The file still exists. It can still be stored, transmitted, and backed up. But its contents are inaccessible without the key. If the server it is stored on is breached, if it is intercepted in transit, if someone gains access to the account it lives in, the encrypted file reveals nothing. It is noise without the key.

The key insight, and the one most people miss, is this: who holds the key matters as much as whether the file is encrypted at all. Many cloud services advertise encryption, and technically they are not lying. Your files are encrypted on their servers. But the service also holds the encryption key, which means that they can decrypt your files whenever they choose to, or are required to. This is called server-side encryption, and while it protects your files from outside attackers, it does not protect them from the company itself, its employees, its legal obligations, or a government that issues a valid request.

What genuinely protects your files is end-to-end encryption with keys that only you hold. In this model, your files are encrypted on your device before they leave it. They travel encrypted. They are stored encrypted. And only your device, using your key, can decrypt them. The service storing your files cannot read them. Neither can a government demanding access. Neither can an attacker who breaches the server. The encrypted data is useless without the key, and the key never leaves your device.

This is the model we built the Vault toward at Cyber Konekt. The principle that your files should be encrypted before they leave your hands, and that nobody except you should ever hold the key to them, is not technically complicated to implement. It is just not the default, because the default serves the interests of the service provider more than the user. Changing it requires deliberate choice.

Practical File Protection: Where to Start

For most people, meaningful file protection involves three practical steps, none of which requires technical expertise. The first is being selective about what goes to cloud storage. Not everything needs to live in the cloud. Sensitive documents, identity papers, medical records, and financial files can be stored locally and backed up to an encrypted external drive. A physical backup is less convenient than cloud storage, but it sits entirely within your control. Nobody can breach a drive that is not connected to the internet.

The second is using a cloud service that offers genuine end-to-end encryption with zero-knowledge architecture, meaning the service itself cannot read your files. Proton Drive and Tresorit are among the established services in this category. They cost money, which is worth framing as the subscription fee for a service whose business model is protecting your files rather than monetising them.

The third is encrypting your most sensitive individual files before storing them anywhere. Tools like VeraCrypt, which is free and open source, allow you to create an encrypted container, essentially a locked vault within your device, into which you can place files you want to keep completely private. Accessing them requires a password. Without the password, they are noise. This layer of protection survives even a device breach, because the encrypted files themselves reveal nothing.

None of these steps is difficult. All three together represent a substantially more secure file posture than the default settings most people never changed.

The Message You Sent Is Not as Private as You Think

Most people sending a message today are using one of a handful of apps: WhatsApp, iMessage, Instagram Direct, Telegram, or one of the platforms' built-in messaging systems. Most of these apps use end-to-end encryption for message content, which means the content of what you type cannot be read in transit or by the service. This is a genuine protection and worth acknowledging.

But there is a dimension of messaging that encryption does not touch. It is called metadata, and it is where the real story lives.

Metadata is data about data. In the context of messaging, it means not what you said, but who you said it to, when, how often, for how long, from where, and using what device. WhatsApp, owned by Meta, is transparent in its privacy policy about collecting this information. It knows who you message and when. It knows your IP address, which reveals your approximate location. It knows your device details, your network information, and identifiers linked to your other Meta accounts. It shares this metadata across Meta's services, where it contributes to the behavioral profile used to target you with advertising.

The former NSA director Michael Hayden has said, in a widely quoted remark, that the agency kills people based on metadata. That is an extreme context, but the underlying point is serious: metadata can be as revealing as content, sometimes more so. Knowing that you called a cancer specialist three

times last week tells a story without requiring anyone to hear what was said.

Signal, and Why It Is Different

Signal reached 70 million users in 2024, a figure that would have seemed implausible a few years earlier for an app whose entire value proposition is what it does not do.

Signal does not collect metadata in the way other messaging apps do. It stores almost nothing: your phone number and the date you last connected to the service. It cannot tell law enforcement who you have been talking to, because it does not know. It cannot tell advertisers anything about your behavior, because it collects nothing to tell them. It is funded by donations and grants rather than data, which means its incentive structure is genuinely aligned with its users rather than in tension with them.

The Signal Protocol, the encryption system underlying Signal's messages, has been independently audited and is widely regarded by cryptographers as the gold standard for secure messaging. Notably, it is the same protocol that WhatsApp uses for its message encryption. The difference is everything that surrounds that protocol: what data is collected, who owns the company, and what the business model is.

Signal's limitations are worth naming honestly. Its user base, while growing, is smaller than WhatsApp's. Asking people to switch to a new messaging app is a social coordination problem as much as a technical one: the app is only useful if the people you want to talk to are also on it. For many people,

switching entirely is not practical. The realistic approach is to use Signal for conversations where privacy matters most, and to be aware of what you are giving up when you use other platforms.

For those who want to go further, Session is worth knowing about. It removes the phone number requirement entirely, so there is no identifying link between your account and your real identity at all. It also uses a decentralised network, which means there is no central server that could be breached or compelled to hand over data. Messages route through a distributed system in a way that makes even the metadata of who is talking to whom extremely difficult to determine. It is a more specialised tool than Signal, and its user base is smaller, but for conversations where the stakes are high enough to warrant it, it represents a meaningfully different level of protection.

The Email Problem

Email deserves a brief separate mention because it is so deeply embedded in daily life and so consistently misunderstood from a privacy perspective.

Standard email is not encrypted. When you send an email using a major free provider, the content of that email is readable by the provider, who uses it to serve you relevant advertising. The email also passes through multiple servers on its way to the recipient, each of which can theoretically read it. It is, in privacy terms, closer to a postcard than a sealed letter. Encrypted email services like Proton Mail, which is based in Switzerland under strong privacy law and uses end-to-end

encryption for messages between Proton users, offer a meaningfully different option. Messages between Proton users are encrypted in a way that Proton itself cannot read. Messages sent to users of other email services are encrypted in transit but readable at the destination, which is a limitation inherent to how email works rather than a flaw in Proton's design.

Switching your primary email to an encrypted service is a bigger lifestyle change than switching a search engine. But for sensitive communications, financial matters, medical correspondence, and anything you would be uncomfortable seeing read aloud in a courtroom, an encrypted alternative is worth considering.

Five Steps to Protect Your Files and Communications

1. Audit your cloud backup settings today. Check which folders and file types are being backed up automatically, and remove sensitive documents, identity papers, and medical files from automatic sync. Moving them to a local encrypted folder instead is worthwhile.
2. For sensitive files that must live in the cloud, use a zero-knowledge encrypted service like Proton Drive or Tresorit. The subscription cost is the price of a service that cannot read your files even under legal pressure.
3. For your most sensitive individual files, use VeraCrypt to create an encrypted container on your device. Free, open source, and used by journalists and researchers who need genuine file security.
4. For private conversations, install Signal and use it for communications where privacy matters. You do not need to abandon other messaging apps entirely. Use the right tool for the right conversation.
5. For sensitive email, consider Proton Mail for correspondence involving health, finances, legal matters, or anything else you consider genuinely private. Free tier available, paid plans offer more storage and features.

The photograph of your passport has been sitting in a cloud server for four years. It is probably fine. Probably is doing a lot of work in that sentence.

CHAPTER 10

DIFFERENT THREATS FOR DIFFERENT LIVES

Daniel is thirty-eight years old. He lives in Amsterdam with his partner, works as a project manager for a mid-sized logistics company, runs three times a week, and has a standing Wednesday evening call with his mother in Lisbon. He is not famous. He is not wealthy. He has done nothing that any reasonable person would describe as interesting from a surveillance perspective.

He is also generating data at a rate that would have been unimaginable twenty years ago. Not because he is careless, or technically unsophisticated, or unusually reckless with his personal information. Simply because he is alive and connected, and that is now enough.

This chapter follows Daniel through a single ordinary day. In each part of it, the data being generated about him is different, the systems collecting it are different, and the potential for harm is different. The same is true of your day. And the day before that.

Daniel at Work

Daniel arrives at his office at 8.47 on a Tuesday morning. He badges in at the entrance, which logs his arrival time to the nearest second. He sits down, opens his laptop, which is company-owned, and begins his day.

The laptop belongs to his employer. So does everything on it. His emails, his calendar, his documents, his browser history, and the messages he sends through the company's internal communication platform are all, legally and practically, company property. Most employment contracts in Europe and North America are explicit about this, somewhere in the clauses most people do not read before signing.

What Daniel does not know, because it was not in the onboarding presentation, is that his company uses monitoring software that logs his keystrokes, takes periodic screenshots of his screen, and tracks which applications he spends time in throughout the day. This is not unusual. According to an ExpressVPN survey of 1,500 US employers and 1,500 employees, 73 percent of companies use online tracking tools to monitor websites, applications, or screen activity. In Europe the regulatory framework is tighter, but monitoring remains widespread and is legal provided employees are informed, which most are, in documents they did not read.

His company also uses an AI-powered productivity analytics tool that generates a weekly score for each employee based on their digital activity patterns: the time between emails sent, the frequency of their system interactions, the ratio of active to idle time on their devices. Daniel's score last week was 73

out of 100. He has never been shown this score. It exists in a dashboard that his line manager reviews monthly. It has already influenced one performance conversation without Daniel knowing it was part of the picture.

When Daniel uses his work phone to message a colleague, that message passes through the company's mobile device management system. When he connects to a client call from home on a Friday afternoon, his home IP address is logged in the company's access records. When he eventually leaves this job, voluntarily or otherwise, several years of his professional behavior will remain in his employer's systems, accessible to whoever replaces his manager and whoever comes after that.

The monitoring does not make Daniel work worse. But it changes how he works. He has noticed, without quite articulating it to himself, that he pauses before opening certain websites during the working day. That he reads certain internal messages more carefully than he used to. That he feels a low-level awareness of being observed that was not there when he started. He is one of the 59 percent of employees who report stress or anxiety associated with workplace surveillance, though he would not describe his feeling as stress. He would describe it as just how it is now.

Daniel at Home

Daniel's apartment has a smart speaker in the kitchen, a connected television in the living room, a video doorbell at the entrance, and a wireless router that his internet provider installed and that he has never logged into. He also has a robotic vacuum cleaner that he bought because it was on sale and which, it turns out, has been mapping his apartment floor plan since the day he set it up.

Ninety-three percent of Americans own at least one smart home device, a figure that holds broadly across Western Europe too. Most people think of these devices as tools that make their homes more convenient. They are also, simultaneously, sensors that continuously generate data about the most intimate space in their lives.

The smart speaker processes audio. Its manufacturer insists it only activates on the wake word, and this is mostly true. But the clips of audio sent to company servers for processing, including occasional accidental activations, are stored, reviewed by human contractors in some cases, and used to improve voice recognition systems. Amazon's Alexa, a 2024 study by Surfshark found, collects 28 out of 32 possible data points, including precise location, contact information, and health data, more than three times the average smart home device. Daniel did not read the privacy policy when he set it up. It was thirty-four pages long.

The video doorbell captures footage of everyone who approaches his building. This footage is stored on a cloud server. In the United States, Ring, which is owned by Amazon,

had a documented history of sharing video footage with over two thousand law enforcement agencies without requiring a warrant, a practice it ended only after significant public and regulatory pressure. Daniel's doorbell is a different brand, but the architecture is similar: footage he captures sits on someone else's server, under someone else's terms.

There is also the security camera Daniel installed eighteen months ago after a neighbour's bike was stolen from the building entrance. It points at his front door from inside the hallway. He chose a budget model, set it up in twenty minutes, and has not thought about it since.

A 2025 study by the cybersecurity firm Bitsight found over 40,000 internet-connected cameras streaming live footage online with no password protection whatsoever, accessible to anyone who knew where to look. The list included home security cameras, baby monitors, and pet cams. No hacking required. A web browser and the right IP address were enough. Some of those feeds were appearing on dark web sites where people sell access to live footage from strangers' homes. And then there is the laptop sitting open on Daniel's kitchen table, its webcam pointing at the room, its indicator light dark. No password issues there. Probably. Daniel's hallway camera almost certainly has a password. Almost certainly.

The connected television watches Daniel as much as Daniel watches it. Automatic content recognition technology, embedded in most modern smart TVs and switched on by default in every major brand as of 2026, takes screenshots of whatever is displayed on screen up to twice every second. It

does not matter whether Daniel is watching a streaming service, a broadcast channel, or a gaming console connected through HDMI.

The TV is capturing all of it, fingerprinting the audio and the image, and transmitting data to the manufacturer's servers roughly every fifteen to sixty seconds, via an accessible internet connection. A peer-reviewed study published in 2024 by researchers at UC Davis, University College London, and Universidad Carlos III de Madrid measured this directly on Samsung and LG televisions by recording the actual network traffic leaving the devices in real time. Neither Samsung nor LG issued a public rebuttal. Within a year, the findings had been cited in legal action by a state attorney general, and Samsung had settled.

Daniel watches a documentary about climate migration late on a Sunday evening. By Monday morning, this information has contributed to an updated profile that will influence the advertisements he sees for the next several weeks. He could turn ACR off in the settings menu. But the same research found that firmware updates silently re-enable privacy settings, often without notification. His living room is not a private space. It is a broadcast studio he did not know he was running.

His home, which feels like his most private space, is generating a continuous stream of behavioral data. The irony is that Daniel specifically thinks of his home as the place where he is not being watched.

Daniel as a Patient

In February, Daniel notices something that concerns him. A mole on his forearm that looks slightly different from how he remembers it. He does not call his doctor immediately. He does what almost everyone does first: he searches.

The searches happen on his personal phone, using the browser that came with it and the search engine set as default. Within the search session, he also visits two medical information websites and reads an article about melanoma warning signs. He closes the browser. He feels better. He books a doctor's appointment for the following week.

What Daniel does not know is that those searches have already been associated with his advertising profile. Health-related search queries are among the most commercially valuable data points in the digital advertising ecosystem. Within 48 hours, Daniel begins seeing advertisements for dermatology clinics, sun protection products, and a health insurance top-up plan. None of this is illegal. None of it required his explicit consent beyond the terms he agreed to when he first used the search engine.

His doctor's appointment goes well. The mole is benign. He is referred to a dermatologist for routine monitoring, which requires him to download a health tracking app that his dermatology clinic recommends for photographing and tracking skin changes over time. The app asks for access to his photo library and his location. He agrees, because the clinic told him to use it and it seems medical.

The app is not covered by HIPAA, because it is a consumer product rather than a clinical tool. Approximately 320 million global health app users sit in this same regulatory grey zone, their health data governed by consumer privacy law and the app's own privacy policy rather than medical data protection frameworks. The app's privacy policy, which Daniel did not read, permits sharing of anonymised data with research partners and advertising networks. The data includes the photographs he uploads, the frequency of his check-ins, and the health notes he types into the journal section.

The appointment outcome was good. But the data trail the appointment generated will persist long after Daniel has stopped thinking about it.

Daniel as a Citizen

Daniel votes. He attends an occasional public demonstration, the kind where several thousand people gather in a city square to express an opinion about something that matters to them. He signed an online petition last year about housing policy. He donates a small monthly amount to an environmental organisation. He follows several political accounts on social media, some of which hold views that are not identical to those of his employer or his landlord.

None of these acts is secret. All of them are legal. All of them generate data.

The demonstration he attended was photographed by news outlets, recorded by CCTV systems operated by the city, and documented by the phones of several thousand attendees who

posted images to social media. Facial recognition technology, which is used by law enforcement agencies in several European countries and extensively across authoritarian states, can place an individual at a specific location and time using nothing more than a clear photograph and a database.

His online petition signature is public. His donation to the environmental organisation is a matter of financial record. His social media follows are visible to the platform, which uses them, alongside his broader behavioral profile, to infer his political orientation, his values, and his likely responses to certain types of content. This inferred political profile is not separated from his commercial profile. It informs the advertising he receives, the content he is shown, and, in countries where political advertising targeting is permitted, the political messages he is served.

Daniel lives in a country with strong democratic institutions and robust rule of law. His civic participation carries no immediate risk. But the data generated by that participation does not stay within the boundaries of the current political moment. It is stored. It accumulates. And the political climate of the country he lives in ten years from now is not guaranteed to be identical to the one he lives in today. History suggests that data collected under one political regime has a habit of becoming useful under the next one.

Daniel does not think about this. Most people do not. It is the kind of concern that feels abstract until it is not.

Daniel in a Private Moment

It is a Thursday night in October. Daniel's partner is away for a week visiting family. The apartment is quiet. He has eaten dinner, done the dishes, and is sitting on the sofa with his phone.

He opens a browser and searches for something he has been thinking about for several months. It is not dramatic. It is not illegal. It is simply private: something he has not discussed with anyone, including his partner, and that he would find acutely uncomfortable if someone he knew were to see it. It might be a health concern he is not ready to name. A question about his relationship. A curiosity he would not easily explain. A search for help with something he is ashamed of.

He searches. He reads. He closes the browser.

The search has been logged. The websites he visited have noted his arrival and his behavior through the usual mechanisms: fingerprint, IP address, tracking scripts, behavioral signals. The content of what he was looking at has contributed to a profile update that will influence what he sees for weeks. He does not know this. He was alone in his apartment. He felt private.

This is the moment that sits at the heart of why privacy matters. Not the surveillance of criminals or dissidents, though that matters too. But this: the quiet, ordinary moment of a man sitting alone on a sofa, looking for something he needed, in the only space he had to look. The feeling of privacy he had in that moment was real to him. The privacy itself was not.

There is nothing he did wrong. There is nothing he should be ashamed of. The shame, if it belongs anywhere, belongs to a system that turned that moment into a data point before he had finished reading the first result.

One Person, Every Threat Model

Daniel is not a high-risk individual. He is not a journalist protecting a source, or an activist in a country where his beliefs could get him imprisoned, or a survivor of abuse trying to stay hidden from someone who wants to find him. He is an ordinary man living an ordinary life in a stable democracy with reasonable privacy laws.

And yet across a single day, he has been monitored at work by systems he was never shown. He has had his home's interior documented by devices he thought of as appliances. He has had a health concern turned into an advertising signal before he had finished processing it himself. His civic participation has been logged in ways that could matter in futures he cannot predict. And a private moment of genuine need was recorded without his awareness.

The concept of a threat model, the security researcher's tool for thinking about who might want your data and what they might do with it, is sometimes presented as something that only matters for people with unusual circumstances. This chapter is an argument against that framing. Daniel's threat model is not one thing. It is five things, running simultaneously, generating data that accumulates across every

context of his life without pause or separation.

The tools in this book are not designed for people with unusual circumstances. They are designed for Daniel. For the 74 percent of employees whose employers are monitoring them in ways they may not fully understand. For the hundreds of millions of people whose health data sits outside meaningful legal protection. For every person who has ever searched for something private and then wondered, briefly, who else might know.

That is most people. It is probably you. It was certainly Daniel, sitting alone on a Thursday night in October, feeling private in a room full of open windows.

Five Context-Specific Steps Worth Taking

1. At work: assume everything on a company device is visible to your employer and act accordingly. For genuinely private communications during working hours, use your personal device on your personal mobile data connection, not the office wifi.
2. At home: audit your smart devices. Check each one's privacy settings, disable microphones and cameras you do not actively use, and consider whether the convenience of each device is worth the data it generates. For sensitive conversations, a room with no smart speakers is still possible to find.
3. As a patient: before downloading any health or wellness app, check whether it is covered by HIPAA or equivalent medical data law in your country. If it is a consumer app, assume your data will be treated as consumer data. Use a private browser for health searches on sensitive topics.
4. As a citizen: if you attend demonstrations or civic events where you would prefer not to be identified, consider leaving your phone at home or in a Faraday bag, which blocks all signals. Your phone's presence at a location is a data point regardless of whether you take it out of your pocket.
5. In private moments: use a privacy-focused browser with tracking protection enabled for searches you consider genuinely personal. The steps from Chapter 8 apply here most directly. A private browsing window is not enough. A browser designed to resist tracking is.

Daniel closed the browser. The search remained open somewhere else.

CHAPTER 11

HOW TO SPOT A PRIVACY TOOL THAT IS LYING TO YOU

There is a particular category of product that I have thought about a great deal since starting Cyber Konekt. It is the privacy tool that is not, in any meaningful sense, a privacy tool. It looks like one. It uses the right words. It has a shield logo and a dark colour scheme and marketing copy that talks about protecting your data and keeping you safe online. It may even have a legitimate technical function underneath all of that.

But somewhere between the marketing and the mechanism, something went wrong. Or more precisely, something was designed to go wrong, in a way that is invisible to the person downloading it but visible immediately to anyone who knows where to look.

I want to show you where to look.

The Free VPN Problem, Stated Plainly

The single most systematically misleading category of privacy product available today is the free VPN with no clear business model. Not all of them. But enough that the category as a whole deserves a level of suspicion that most users do not apply to it.

A VPN costs money to operate. Servers, bandwidth, maintenance, support: none of these are free. When a VPN is offered at no charge, the question worth asking is not whether to be grateful, but how the company is covering its costs. Because it is covering them somehow. And in the privacy tool space, the answer to that question determines whether you have a protection or a trap.

A 2024 study found that 88 percent of free Android VPNs leak user data. Not occasionally. Systematically. These are not fringe applications downloaded from obscure websites. Many of them have millions of downloads from official app stores, four-star ratings, and professional-looking landing pages that talk about privacy in fluent, confident terms.

The mechanism is straightforward. A free VPN encrypts your traffic, which is real. But many also log your browsing activity, your IP address, your device identifiers, and your behavioral patterns. And then sell that data to advertising networks and data brokers. You came to the product to escape surveillance. That product is the surveillance, wearing a different coat.

This is not a design flaw. It is a business model. And it works because most users never think to ask how a free service covering real infrastructure costs is managing to stay in business.

The Language That Should Stop You Cold

Privacy marketing has developed a vocabulary that sounds protective and means almost nothing. Learning to read it is one of the most useful skills available to anyone trying to evaluate a privacy product honestly.

The phrase that does the most damage, buried deepest, and sounds the most innocuous is this: 'we may share your data with trusted partners.'

It appears, in some form, in the privacy policies of a remarkable number of products that lead with strong privacy claims in their marketing. The opening of the policy talks about protecting your data, respecting your privacy, never selling your information. Then, twenty paragraphs later, in a section titled 'Disclosure of Information' or 'Third Party Relationships' or something equally bureaucratic, the clause appears. Trusted partners. A phrase with no legal definition, no limit on who qualifies, and no requirement to tell you who they are or what they do with what they receive.

'Trusted' in this context means trusted by the company. Not by you. The company trusts the advertising network that pays it for access to aggregated behavioral data. The company trusts the analytics platform that processes your usage patterns. The company trusts the data broker that buys monthly exports of user profiles. All of these are trusted partners. None of them are who you thought of when you read the word privacy on the download page.

The other phrase worth treating with immediate suspicion is 'analytics and product improvement.' This is the carve-out that permits profiling while sounding like responsible engineering. Analytics requires data. Product improvement requires data. Both are true. But the scope of what is collected under this justification is frequently far broader than what any honest reading of product improvement would require, and the data collected is often retained indefinitely, shared with third parties, and used for purposes that have nothing to do with making the product better.

Read privacy policies from the middle, not the beginning. The opening is marketing. The carve-outs are near the end.

What the Business Model Tells You Before You Read Anything Else

There is a test you can apply to any privacy product before reading a single line of its privacy policy, and it is more reliable than most things that come after.

Ask: how does this company make money?

If the answer is immediately clear, a subscription fee, a one-time purchase, an open-source model funded by donations, that is a meaningful signal. The company's revenue does not depend on your data. Its incentives are aligned with protecting you rather than monetising you. This does not guarantee the product is good. But it removes the most obvious structural conflict of interest.

If the answer is unclear, if the product is free and the revenue model is not explained, if the privacy policy refers to advertising partnerships or data sharing without specifying the scope, if the company is owned by a larger entity whose primary business has nothing to do with privacy: these are signals worth taking seriously. Not as proof of wrongdoing, but as reason to look harder before trusting.

The privacy industry is not immune to the dynamics described in Chapter 2. A company that pivots to privacy branding because it is commercially attractive, without restructuring the underlying data architecture, is not a privacy company. It is a company with a privacy brand.

The difference matters enormously and is almost never visible on the surface.

The Permissions That Give It Away

For mobile apps specifically, there is another layer of evaluation available that requires no technical knowledge. It is the permissions screen.

When you install an app, it requests access to parts of your device: your camera, your microphone, your location, your contacts, your storage. These requests are visible before you complete the installation. Most people tap through them without reading. But for a privacy tool, they are one of the clearest signals available.

A VPN needs network access. It does not need your contacts. A private browser needs storage access. It does not need your precise location at all times. A secure notes app needs storage. It does not need your camera. When a privacy-branded app requests permissions that have no obvious relationship to its stated function, the most charitable interpretation is poor engineering. The less charitable interpretation is that the permissions are the function, and the stated purpose is the cover.

Cross-reference the permissions against the privacy policy. If the app requests access to your contacts and the privacy policy mentions that contact data may be shared with partners for service improvement, you now have a complete picture of what is happening. The marketing said privacy. The permissions said contacts. The privacy policy said partners.

The actual product is a contact list harvester with a VPN attached.

A Question I Have Had to Answer Myself

I want to be honest about something that makes this chapter more than an exercise in spotting other people's bad behaviour.

Building a privacy-first product under commercial pressure is not a theoretical exercise. It is a series of decisions made under real constraints, with real consequences, over a sustained period of time. And the pressure does not always come from obviously bad actors or obviously bad intentions. It comes from the ordinary logic of building a company.

When we were developing VPN Toolkit, the question of analytics came up early and came up repeatedly. Every product team that is serious about building something good wants to understand how users actually use it. Where they drop off. What features they engage with. What is confusing. What is working. This is not an unreasonable thing to want. In fact, it is how good products get built.

The standard tools for gathering this understanding are analytics platforms that sit inside the app and report back on user behaviour in detail. Firebase. Mixpanel. Amplitude. These are what most apps use. They are genuinely useful. They also require collecting more data about users than we were prepared to collect, routing it through third-party systems, and accepting that the insight we gained came at a cost to the people using the product. For an app that exists to protect its

users from exactly that kind of data collection, the contradiction was not something we could argue our way around.

So we declined standard analytics integrations. Not once, as a single principled moment, but repeatedly, as an ongoing architectural position. There was no dramatic meeting. No founder speech. Just a recurring conversation that kept arriving at the same conclusion: we cannot build a privacy tool on an infrastructure that undermines the thing we are claiming to protect.

What that decision cost us is straightforward. We have less granular insight into how users move through the product than we would otherwise have. We make product decisions with less data than our competitors. We have had conversations with investors and advisors who consider standard analytics a basic hygiene requirement for any serious product, and we have had to explain, more than once, why we see it differently. The explanations are not always received as the obvious position.

The position eventually led us somewhere we did not initially expect. We kept asking the same question: is there a tool that gives us meaningful product insight without requiring us to collect personal data? For a long time the honest answer was no. Then we found TelemetryDeck.

TelemetryDeck is a German analytics platform built from the ground up around genuine anonymisation. It does not collect personal data. It does not store identifiable user information.

User identifiers are first hashed and salted on the device before they leave it, then hashed again with TelemetryDeck's own salt on their servers, and the identification date is never stored, meaning re-identification is practically impossible given current technology even with additional data. Under the GDPR's own definition, based on current technology and GDPR interpretation, the data TelemetryDeck processes does not qualify as personal data at all. It is not privacy-respecting analytics in the sense of trying hard not to collect too much. It is analytics whose architecture makes personal data collection practically impossible given the current state of technology.

This is the distinction the rest of this chapter is built around. Not whether a tool claims privacy, but whether its architecture genuinely delivers it. TelemetryDeck cleared that bar in a way that Firebase, Mixpanel, and Amplitude do not. We are integrating it into VPN Toolkit. We checked first.

The website is a different context again. A public website serves a different purpose: it helps people find us, understand what we do, and decide whether to download the product. Understanding broadly how that works is a legitimate operational need. For this we use Plausible Analytics, a bootstrapped, subscription-funded company that collects no personal data, uses no cookies, and builds no individual user profiles. It is, in the language of this chapter, a company whose business model and whose product are in alignment. You can verify this by reading their policy, which is notably short.

I tell this story not because it makes Cyber Konekt exceptional. Many people building in this space make similar decisions, often under greater pressure. I tell it because I think

it illustrates something the rest of this chapter has been building toward: the difference between a company that treats privacy as a principle it has to live up to under pressure, and one that treats it as a position it holds until the pressure appears.

That difference is not always visible from the outside. But there are ways to look for it.

Five Things Worth Checking Before You Trust Any Privacy Tool

None of these require technical expertise. They require only the willingness to look for a few minutes before committing your data to something.

First, find the business model. If the product is free, find out how the company pays its bills. If the answer involves advertising, data partnerships, or is simply unclear, treat that as a reason to look harder.

Second, read the middle of the privacy policy. Search the document for the phrases 'trusted partners,' 'third parties,' 'analytics,' and 'product improvement.' Read the sentences around each one. The carve-outs live there.

Third, check the permissions on install. For mobile apps, compare what the app asks for against what it needs to do its job. Permissions that do not match the function are a signal. Fourth, look for independent audits. A privacy tool that has had its claims independently verified by a reputable third

party is making a meaningfully different commitment than one that simply says trust us. Audits are not cheap or easy. Companies that commission them are demonstrating something about their priorities.

Fifth, check who owns it. Several privacy-branded products are owned by companies whose primary business is advertising or data collection. Parent company ownership is usually disclosed in the app store listing or the privacy policy. It is worth knowing.

The Five-Minute Privacy Tool Test

1. Search '[product name] business model' or '[product name] revenue.' If the answer is unclear or involves advertising, proceed with caution.
2. Open the privacy policy and search for 'trusted partners,' 'third parties,' and 'analytics.' Read those sections carefully, not the introduction.
3. On mobile, check the permissions the app requests before installing. Compare them against what the product actually needs to function.
4. Search '[product name] audit' or '[product name] independent review.' If a reputable third party has verified their privacy claims, that matters.
5. Search '[product name] parent company' or '[product name] owned by.' Several privacy tools are subsidiaries of data collection companies. This information is rarely prominent.

The shield logo means nothing. The business model means everything.

PART 4
WHAT COMES NEXT

CHAPTER 12

WHAT GOOD REGULATION LOOKS LIKE

In May 2018, a week after the GDPR came into force, a nine-year-old Austrian boy named Max filed a data access request with a gaming company that had been collecting his personal information. He was nine. He had a right, in legal terms equivalent to any adult in the European Union, to know what data the company held about him, on what basis, and who had seen it.

Whether he actually filed that request, or whether a parent did it on his behalf, matters less than the principle it illustrates. The GDPR did not just regulate companies. It created rights. Rights that belong to every person in the EU regardless of age, income, technical knowledge, or whether they can afford a lawyer. That shift, from a system where data protection was something companies promised to a system where it was something individuals could legally demand, is the most important thing the GDPR did.

It is also, in many ways, the floor of what good regulation looks like. Not the ceiling.

This chapter is not a repeat of Chapter 7, which covered the GDPR's enforcement failures in detail. It starts from a different question. Given everything we know about what regulation has failed to achieve, what does success actually require? What does a privacy framework look like when it is working, and where in the world can we see it happening?

What the GDPR Got Right

It is worth being clear about this before moving to what comes next. The GDPR was a genuine shift in the global privacy landscape, and dismissing it because its enforcement has been inconsistent would be as wrong as treating it as a solved problem.

The regulation established that data protection is a fundamental right, not a consumer preference. It moved the legal burden from individuals to organisations, requiring companies to justify their data collection rather than requiring individuals to opt out of it. It created a baseline of rights, access, correction, deletion, portability, and the right to object to automated decision-making, that simply did not exist for most people before 2018. And through the Brussels Effect described in Chapter 3, it exported those rights informally to hundreds of millions of people around the world who had no legal entitlement to them.

These are real achievements. The GDPR also demonstrated something that the privacy debate had previously treated as uncertain: that comprehensive privacy regulation does not destroy the digital economy. The EU's technology sector did not collapse after 2018. Companies adapted. New compliance industries emerged. The fear that strong privacy rules were incompatible with innovation turned out to be, in most respects, unfounded.

But the GDPR was written for the internet of 2016, and the internet of 2026 has moved considerably. Three gaps stand out above the others.

First, the consent mechanism. The GDPR's requirement for informed, specific consent was correct in principle but has been degraded in practice into the cookie banner, a mechanism so universally resented and so easily manipulated by dark patterns, pre-ticked boxes, difficult-to-find decline buttons, and consent fatigue, that it has arguably made people less thoughtful about privacy rather than more. The mechanism was right. The implementation has been captured by the industry it was meant to constrain.

Second, the enforcement asymmetry described at length in Chapter 7: fines levied but not collected, regulators under-resourced relative to the companies they oversee, and a one-stop-shop mechanism that concentrated the regulation of the world's largest technology companies in a single small national authority. The rules were strong. The institutional capacity to enforce them was not.

Third, and perhaps most consequentially for the future: the GDPR was not written for artificial intelligence. It does not adequately address systems that infer sensitive attributes from innocuous data, generate synthetic personal information, or make consequential decisions about people through processes that no human being fully understands. These are not edge cases. They are the most important privacy challenges of the next decade.

The EU AI Act: A Second Attempt at the Brussels Effect

In August 2024, the European Union's Artificial Intelligence Act came into force. It is the world's first comprehensive legal framework for regulating AI, and it represents something significant: an attempt to govern not just what data is collected about people, but what systems are built with that data and what decisions those systems are permitted to make.

The AI Act takes a risk-based approach. It categorises AI systems by the potential harm they can cause and imposes requirements accordingly. At the top of the hierarchy are systems considered to pose unacceptable risks, which are banned outright. These include AI systems that manipulate people through subliminal techniques, systems that exploit vulnerabilities to distort behavior, and, with limited exceptions, real-time biometric surveillance in public spaces. Below that sit high-risk systems, covering areas like employment screening, credit scoring, critical infrastructure, and law enforcement applications, which face strict requirements for transparency, accuracy, human oversight, and documentation.

The AI Act's relationship with privacy is layered. It does not replace the GDPR. It sits alongside it, applying safety and accountability requirements to AI systems the way product safety law applies to physical goods. Where the GDPR asks who has your data and on what basis, the AI Act asks what systems are making decisions about you and whether those systems are safe.

Together they are intended to form a coherent framework for governing the digital environment that most people now live inside.

Whether the AI Act will achieve its own Brussels Effect, whether companies operating globally will adopt its standards everywhere rather than maintain separate practices for the EU, is not yet known. The legal scholar Anu Bradford, who coined the term for the GDPR's global spread, has been cautiously optimistic but noted that AI regulation involves more direct competition with geopolitical interests than data protection did. The United States and China both have strong commercial and strategic reasons to resist European AI governance standards. This is a different contest from cookie banners.

What the AI Act does demonstrate, regardless of its global reach, is that regulation can evolve to meet new technologies rather than always lagging behind them. The gap between the GDPR's 2016 drafting and the capabilities it was struggling to govern by 2024 was roughly eight years. The AI Act was drafted in awareness of that gap. Whether eight years becomes four, or two, or whether regulation can eventually anticipate rather than react, is one of the defining questions of the next decade.

South Korea: What Enforcement With Real Teeth Looks Like

If the GDPR is the most influential privacy framework in the world, South Korea's Personal Information Protection Act is arguably now the most aggressively enforced.

And the contrast between the two reveals something important about what makes regulation effective in practice rather than just on paper.

South Korea's PIPA was comprehensively overhauled in 2023, transforming what had been a moderately strong framework into one with substantially more power. The Personal Information Protection Commission was elevated to an independent top-tier authority, consolidating enforcement powers that had previously been fragmented across multiple agencies. Fines were restructured to be calculated against total global revenue rather than domestic revenue, closing the loophole that had allowed large international companies to pay penalties that bore no meaningful relationship to their actual scale.

The results were immediate. In 2022, the PIPC fined two major US-based platforms a combined 72 million US dollars for collecting and analysing behavioral data for targeted advertising without proper consent. Appeals were unsuccessful. In 2022, the PIPC fined Google approximately 50 million dollars and Meta approximately 22 million dollars for consent failures. The enforcement record is consistent, technically sophisticated, and demonstrably unintimidated by the scale of the companies being regulated.

In early 2026, South Korea went further still, passing amendments that raised the maximum penalty ceiling to ten percent of total global turnover and introduced personal supervisory liability for chief executives, with these provisions scheduled to take effect from September 2026 as of the date of

writing. This last element is worth dwelling on. Most privacy fines, however large, are treated by technology companies as a cost of doing business: a line item absorbed into compliance budgets. Personal liability for senior executives is a different kind of deterrent entirely. It changes the internal conversation about privacy from a legal and compliance question to a leadership one.

South Korea's approach illustrates a principle that good regulation consistently demonstrates: rules without credible consequences are suggestions. The strength of a privacy framework is not measured by the rights it grants on paper but by what happens when those rights are violated.

Australia: Privacy as a Private Right

Australia's Privacy Act amendments, passed in December 2024, introduced something that most privacy frameworks lack and that advocates have argued for since the GDPR came into force: a private right of action.

Under most privacy frameworks, including the GDPR, enforcement is a matter for regulators. If a company misuses your data, you can complain to a supervisory authority, who may or may not investigate, who may or may not take action, who may or may not collect any penalty. The individual whose rights were violated has no direct legal recourse against the company that violated them. The system depends entirely on regulators having the will, the resources, and the capacity to act.

Australia's new statutory tort for serious privacy invasion changes this. For the first time in Australian law, individuals can sue companies directly for serious breaches of their privacy, without needing to route their complaint through a government body first. Maximum penalties for serious breaches already stood at fifty million Australian dollars or thirty percent of turnover, whichever is higher, a ceiling that meaningfully exceeds the GDPR's four percent of global turnover for the most serious cases.

The private right of action is controversial. The technology industry's objection, that it will generate frivolous litigation and create compliance uncertainty, is not entirely without merit. Class action litigation can be blunt, expensive, and slow. But the counter-argument is equally valid: a regulatory system that places the entire burden of enforcement on under-resourced government agencies is one that can be managed, delayed, and defunded. A system where every affected individual is a potential plaintiff is substantially harder to manage away.

The Australian model is too new to have produced a body of case law that demonstrates whether it works in practice. But it represents a genuinely different theory of how privacy rights get enforced, one that distributes enforcement power rather than concentrating it, and one that could influence how other jurisdictions think about closing the gap between rights on paper and consequences for violating them.

What Good Regulation Actually Requires

Looking across the frameworks described in this chapter, and the failures described in Chapter 7, a picture of what effective privacy regulation requires begins to emerge. It is not a single template. Different jurisdictions have different legal traditions, different institutional capacities, and different political environments. But the components of success are recognisable across all of them.

Rights must be real, not performative. The GDPR's rights are legally meaningful. The consent mechanisms designed to implement them have been undermined by dark patterns and interface manipulation. Good regulation pays attention to how rights are exercised in practice, not just whether they exist in legislation. This means regulating not just the collection of data but the design of the systems through which consent is given and rights are invoked.

Enforcement must be adequately resourced and institutionally independent. South Korea's elevation of its privacy regulator to an independent top-tier authority was not a cosmetic change. It was a structural one that removed the institutional pressures that had previously softened enforcement. Regulators that are dependent on the political goodwill of governments that also court large technology companies will, predictably, regulate less aggressively than their rules permit.

Penalties must bear a meaningful relationship to the scale of the violator. A fine calculated as a percentage of domestic revenue from a single market is not a deterrent for a company whose revenues are global. The shift to global turnover as the

penalty basis, seen in both South Korea's 2023 amendments and the GDPR's original design, is correct. The question is whether fines are actually collected, which brings us back to enforcement. Regulation must evolve faster than technology. The eight-year lag between the GDPR's drafting and the AI capabilities it was struggling to govern is too long. The EU AI Act represents an attempt to reduce this lag. Whether it succeeds depends partly on whether it can be amended and updated as AI capabilities develop, or whether it will become a fixed framework that the technology outgrows the way the GDPR did.

And finally, regulation must account for what happens when it fails. The Australian private right of action is one answer: distribute enforcement power so that no single point of failure can render the system inert. Independent audit requirements, mandatory breach notification, and personal liability for senior executives are others. Good regulation assumes that companies will sometimes violate the rules and builds systems that can respond effectively when they do.

The Progress Is Real, and So Is the Distance

I want to be honest about what this chapter has and has not argued. It has not argued that privacy regulation is working. Chapter 7 made clear that, measured against the scale of the surveillance economy it is trying to govern, current regulation is insufficient. The gap between what the rules say and what the industry does remains large.

What this chapter has argued is that the direction of travel is correct, that examples of effective regulation exist and can be

learned from, and that the tools for doing this better are understood even where they are not yet consistently applied. That is a more modest claim, but it is also a more honest one, and it matters because despairing of regulation entirely is precisely what the surveillance economy's beneficiaries would prefer.

The GDPR changed corporate behavior globally in ways that were not predicted. The AI Act may do the same for algorithmic decision-making. South Korea's enforcement record demonstrates that even the largest technology companies respond to credible legal consequences. Australia's private right of action opens a new front that does not depend on government capacity. These are not solutions. They are signals of what is possible.

The distance between those signals and a world in which most people's data is genuinely protected is still measured in years of legislative effort, institutional investment, and political will. None of that happens automatically. It happens because enough people understand what is at stake and act on that understanding.

Which is, in the end, why books like this one exist.

Five Regulatory Developments Worth Following

1. The EU AI Act's full enforcement timeline runs through 2026. How national regulators apply it to real AI systems, particularly in advertising, employment, and credit, will

determine whether it achieves its intended effect or becomes another framework that looks strong on paper.

2. The United States federal privacy law debate has not ended. The American Privacy Rights Act introduced in 2024 stalled, but the underlying pressure for a federal floor has not gone away. State laws continue to proliferate and diverge, which creates its own pressure for federal harmonisation.
3. South Korea's CEO liability provisions, effective September 2026, will be the first major test of whether personal executive accountability changes corporate privacy behaviour at the leadership level. The outcome is worth watching regardless of where you live.
4. Australia's private right of action for privacy violations is the newest experiment in distributing enforcement power to individuals. Its first significant cases will reveal whether the mechanism works in practice and may influence similar proposals in other jurisdictions.
5. The question of AI and data protection is unresolved globally. How regulators in the EU, UK, US, and Asia apply existing privacy frameworks to large language models and generative AI systems trained on personal data is one of the most consequential open questions in technology policy.

Max was nine years old when he had the legal right to ask a company what it knew about him. Most people in the world still do not have that right. The gap between those two facts is where the work is.

CHAPTER 13

**THE COMPANIES
BUILDING
DIFFERENTLY**

There is a version of this chapter that would be straightforward to write and almost entirely useless. It would list a set of well-regarded privacy tools, describe their features, and conclude that privacy-first technology is flourishing. It would be accurate in its details and misleading in its overall impression.

The reality is that these companies, building genuinely privacy-first products are operating in a market where the dominant economic model is their direct opponent. Every day they compete against free services subsidised by the data of the people using them. Every day they ask users to pay for something that the surveillance economy offers at no charge. Every day they make architectural decisions that cost them insight, growth, and revenue, in service of a principle that most of their competitors treat as a marketing position rather than a design constraint.

Understanding that context is what makes the ecosystem described in this chapter interesting. These are not companies that found an easy niche. They are companies that chose a hard one, and that have built something real inside it. That distinction matters when you are deciding which tools to trust and which companies to support.

The Difference Between Privacy Branding and Privacy Architecture

Chapter 11 described how to spot a privacy tool that is misleading users. This chapter approaches the same question from the other direction: what does a genuine privacy-first company actually look like underneath the marketing?

The honest answer is that even the best ones are imperfect. Signal, which is as close to a gold standard for private messaging as the consumer market has produced, still requires a phone number to register. That phone number is an identifying link between your Signal account and your real identity. Signal publishes this limitation openly in its documentation and has explained the technical and abuse-prevention reasons for it. The transparency is itself a signal of something genuine: a company that acknowledges its constraints rather than hiding them is operating differently from one that presents perfection it has not achieved.

Proton, which has grown from a two-person encrypted email project founded by CERN scientists in 2014 into a suite of privacy tools used by over 100 million people, has faced its own moments of tension. In 2021 it was ordered by Swiss authorities to hand over the IP address of a user under investigation for a protest-related offence. Proton complied, because Swiss law required it, and because Proton Mail had stated in its documentation that it could be compelled to do so under Swiss legal process.

Critics pointed to this as a failure. Proton pointed to it as evidence that their privacy policy was accurate rather than aspirational. Both characterisations are true.

The lesson from both of these cases is not that Signal and Proton are untrustworthy. It is that no company operates in a legal and technical vacuum, and that the question worth asking about any privacy tool is not whether it is perfect but whether it is honest about where it falls short. Companies that publish their limitations, commission independent audits, and respond to criticism by improving their documentation and architecture are demonstrating something that matters: they take the commitment seriously enough to be accountable to it.

Signal: A Different Theory of What a Company Is For

Signal is an unusual organisation by almost any measure. It is a nonprofit funded by donations and grants, with no advertising revenue, no data monetisation, and no investors whose returns depend on growth metrics. Its founder, Moxie Marlinspike, spent years building cryptographic tools for people operating in genuinely dangerous environments before Signal became a consumer product.

The Signal Protocol, the encryption system underlying its messages, is now used by more than a billion people, including through WhatsApp, which adopted it in 2016.

What makes Signal worth discussing in a chapter about building differently is not primarily its technical architecture, which is documented extensively elsewhere, but its organisational architecture. Signal exists to provide a specific service to its users. It does not exist to grow, to capture market share, to achieve a liquidity event, or to generate returns for shareholders. These are not constraints placed on a commercial company by regulation. They are the founding premise of the organisation. The absence of a commercial growth imperative changes almost every design decision that follows from it.

Signal reached 70 million users in 2024 without a single dollar of advertising spend and without selling or sharing a single piece of user data. That is not a business achievement. It is a proof of concept for a different theory of what technology organisations can be built to do.

Proton: Structural Protection of Mission

Proton's response to the tension between commercial sustainability and privacy commitment has been architectural in a different sense: not technical architecture but organisational. In June 2024, Proton transitioned to a nonprofit foundation structure, with the newly created Proton Foundation becoming the company's primary shareholder.

The significance of this move is structural rather than symbolic. The Proton Foundation's mandate is to protect the company's founding mission. As primary shareholder, it has the voting power to block any acquisition, merger, or strategic pivot that would compromise that mission. A future

leadership team that decided to introduce advertising, sell user data, or pivot away from privacy could not do so without the foundation's consent. The privacy commitment is not held in place by the values of the current team alone. It is held in place by the ownership structure.

This matters because the history of privacy-focused companies includes several cases where the original commitment survived exactly as long as the original founders remained in control. Companies get acquired. Teams change. Investors' priorities shift. Proton's foundation structure is an attempt to make the mission durable in a way that personal commitment alone cannot guarantee. Whether it succeeds over the long term depends on factors that cannot be known in advance. But it represents a more serious attempt to solve the problem of mission persistence than most companies have made.

Note: NordVPN is a Cyber Konekt affiliate partner. Full disclosure in the author's note.

NordVPN: Verification as a Standard

NordVPN occupies an interesting position in the privacy ecosystem. It is one of the most commercially successful VPN providers in the world, with millions of users and a significant marketing budget. It is also, as of late 2025, the only major VPN provider to have had its no-logging policy independently verified six times by Deloitte, one of the largest auditing firms in the world.

The audit record matters in the context of Chapter 11's argument about how to evaluate privacy tools. Most VPN providers claim a no-logging policy. NordVPN has subjected its claim to independent verification annually since 2018, using a consistent auditing standard, the ISAE 3000 framework, and making the reports available to users. This is not proof of perfection. Audits verify what auditors can inspect at the time of inspection. But it is a meaningfully higher standard of accountability than a policy document that nobody has checked.

The broader point NordVPN illustrates is that commercial success and genuine privacy commitment are not inherently incompatible. The argument that only nonprofit or donation-funded organisations can be trusted with privacy is too simple. What matters is whether the business model creates incentives that align with or contradict the privacy commitment, and whether the company submits its claims to external verification rather than asking for trust on the basis of marketing alone.

The Less Visible Infrastructure

The companies most people have heard of represent the visible layer of the privacy ecosystem. Below them sits an infrastructure of less prominent but equally important organisations whose work makes privacy-first technology possible at scale.

Let's Encrypt is a nonprofit certificate authority that provides free SSL certificates to websites, enabling encrypted connections between websites and their visitors. Before Let's

Encrypt launched in 2016, implementing HTTPS on a website required paying for a certificate, a barrier that kept a significant portion of the web unencrypted. Today, more than 400 million websites use Let's Encrypt certificates. The encrypted web, which most users now experience as a default, was substantially built by a nonprofit giving away a service for free because it believed encryption should not be a premium feature.

The Tor Project, another nonprofit, maintains the network that underpins the most widely used anonymity tool in the world. The Tor Browser routes internet traffic through a series of volunteer-operated servers, obscuring the user's location and identity in a way that commercial VPNs cannot match. Tor is used by journalists, activists, and dissidents in countries where the consequences of being identified are severe. It is also used by ordinary people who simply want to browse without being profiled. The organisation that maintains it operates on grants and donations, without commercial revenue, because the population most in need of its service is not in a position to pay for it.

The Electronic Frontier Foundation has been litigating for digital rights since 1990, long before most of the companies described in this chapter existed. It has fought cases that established the legal frameworks within which privacy tools operate, and it publishes detailed technical guides that have shaped the thinking of a generation of privacy engineers and advocates. It is not a technology company in any conventional sense, but the ecosystem described in this chapter would look significantly different without it.

These organisations do work that commercial companies cannot or will not do, because the work does not generate revenue. They are the foundation layer of the privacy ecosystem, and the fact that most people have never heard of them is part of why the visible layer of consumer privacy tools is able to exist and function.

Where Cyber Konekt Sits in This

I want to place Cyber Konekt in this ecosystem honestly, which means neither understating what we are trying to do nor overstating where we are.

We are a small team based in Prague, building privacy tools for everyday users who should not need a computer science degree to protect themselves online. The founding conviction behind Cyber Konekt is the same conviction that runs through this book: that privacy is not a technical problem to be solved by experts for experts, but a basic right that should be accessible to anyone who wants it. The tools we build are designed around that conviction.

VPN Toolkit, our flagship Android app, bundles the layered defense described in Chapter 8 into a single place, because the evidence we have seen consistently is that complexity is the main obstacle between people who care about privacy and people who act on it. The file Vault we are building extends the same principle to personal files, with an architecture designed around user-held keys and built on a rigorously audited cryptographic foundation. The cloud backup service in development is designed on the same architecture: your data, your keys, our infrastructure, no access.

We are pre-revenue on several of these products and actively building toward a future where the privacy-first toolkit we have assembled is something an everyday user can pick up and use without reading a technical guide first. We have made the kinds of decisions described in Chapter 11: making deliberate decisions about which analytics tools meet our privacy standard and which do not, choosing tools whose architecture genuinely rules out re-identification over standard integrations that cannot make that claim, choosing infrastructure partners whose own privacy commitments we can verify, and treating every architectural decision as a question about whose interests the system serves.

We are not Signal. We are not Proton. We are a company at an earlier stage of a longer journey, building in a space that the organisations described in this chapter have helped make possible. The best thing I can say about where we sit in this ecosystem is that we are trying to earn our place in it by the same standard we have applied to everyone else in this chapter: not by claiming privacy, but by building it.

What This Ecosystem Needs From You

The companies and organisations described in this chapter share a structural disadvantage relative to the surveillance economy they are trying to displace. They cannot outspend it. They cannot match its scale. They cannot compete on the dimension of free, because their model depends on not monetising users.

What they can compete on is trust, and trust is built by users choosing them. Every subscription to a privacy-first service is a vote for a different model. Every person who switches from a surveillance-funded product to a privacy-funded one shifts, by a small but real amount, the economic viability of building differently. The privacy ecosystem does not grow because of regulation alone, or because of technology alone. It grows because enough people decide that the price of free is too high, and act on that decision.

That decision does not have to be absolute or immediate. Nobody is asking you to abandon every product you currently use on the same afternoon. What the ecosystem asks for is something more modest: a direction of travel. Each time you have a choice between a tool that treats your data as its product and one that does not, choose the one that does not. Over time, those choices add up to something.

The Ecosystem at a Glance

1. Signal: Nonprofit messaging app, no advertising, no data monetisation, 70 million users. The gold standard for private communication, with honestly documented limitations.
2. Proton: Swiss-based privacy suite covering email, VPN, cloud storage, and more. 100 million users. Transitioned to nonprofit foundation ownership in 2024 to make the privacy mission structurally durable.
3. NordVPN: Commercially successful VPN provider with six independent Deloitte audits of its no-logging policy since 2018. Demonstrates that commercial scale and genuine privacy commitment can coexist.
4. Let's Encrypt: Nonprofit certificate authority that made encrypted web connections free and universal. Over 400 million websites. The infrastructure layer most users never see.
5. The Tor Project: Nonprofit maintaining the world's most robust anonymity network, used by journalists, activists, and ordinary people who want to browse without being profiled. Funded by donations and grants.

None of these companies are perfect. All of them are trying. In this industry, at this moment, that distinction is everything.

CHAPTER 14

YOU ARE NOT POWERLESS

I want to tell you something honest before this book ends.

When I look at the full scale of what we are up against, not just the surveillance capitalism infrastructure described in these pages but the degree to which governments around the world have aligned themselves with that infrastructure, something in me hesitates. Not occasionally. Regularly. The size of the system, the depth of the economic interests defending it, the speed at which it evolves, the political will required to constrain it: these are not small things. They are genuinely large things, and anyone who tells you otherwise is either not paying attention or trying to sell you something.

I, along with my team, have built a company in this space. We have made the decisions described in this book: the harder architectural choices, the declined partnerships, the slower growth. I have had the conversations with people who thought those decisions were naive. And I have felt, more mornings than I would like to admit, the weight of the question underneath all of this. Does it matter? Is an individual standing up to a system this large, saying I want things to be better and I am going to do something about it, a meaningful act or a gesture?

The doubt returns. That is the honest answer. It returns because the system is real and it is large and it does not particularly care about any individual's objection to it. But something else returns with it, and this part has grown stronger over time, not weaker. A clearer understanding of what I am actually doing and why. A recognition that the enemies of privacy and freedom do not simply hand over

better outcomes. That these are hard-won in specific arenas by specific people who decided to show up and keep showing up. And an understanding that what I am doing is not primarily for myself. It is a contribution to something I want to pass on, a world in which the default setting is not surveillance, in which ordinary people have real control over their own lives, in which the tools to protect yourself are accessible rather than reserved for those with technical knowledge or the money to pay for premium privacy.

The confidence to hold that position has evolved. That is the right word for it. Not arrived. Evolved. Each day of continuing, each decision made in the right direction, each person who downloads VPN Toolkit and understands something they did not understand before: these things accumulate. The doubt does not disappear. But the person answering it becomes more capable of carrying it.

I am telling you this because I think it applies to you too. Not in building a company, necessarily. But in deciding whether any of what this book has described is worth acting on. The answer to that question is not found once and settled. It is found again each day, and it gets easier to find each time you have found it before.

What You Have Actually Done by Reading This

There is a tendency in books like this one to understate the significance of the reader's act of reading it. To treat the knowledge just transferred as merely preparatory, a precondition for the real work that comes next. I do not think that is quite right.

Chapter 4 described the chilling effect: the way that awareness of surveillance changes behavior, suppresses curiosity, narrows the space for free thought. The mechanism works because most people do not understand what is happening to them or why. Ignorance is not a passive state in the surveillance economy. It is an asset, extracted and monetised daily from billions of people who have no idea the transaction is taking place.

You now understand the transaction. That understanding is not neutral. It changes how you see a cookie banner, a free app, a smart speaker, a privacy policy. It changes what questions you ask before you agree to something. It changes the conversations you have with people around you. Awareness, passed from one person to another, is how the culture around privacy changes. Not through legislation alone, not through technology alone, but through the slow accumulation of people who know what they are consenting to and choose differently because of it.

This book is already part of that. So are you, now, in a way you were not before you opened it.

The Things That Are Changing

Thirteen chapters of honest analysis of a serious problem can leave a reader with a specific kind of heaviness if the forward-looking section does not do its job. I want to do that job properly, which means not telling you that everything will be fine, because I do not know that, and the earlier chapters have not earned that claim. What I can tell you, honestly, is what is actually changing.

Regulation is moving. Slowly, unevenly, with the enforcement gaps described in Chapter 7 still very much present. But the direction is consistent. More countries have comprehensive privacy legislation today than at any point in history. The EU AI Act represents an attempt to govern algorithmic systems that the GDPR was not equipped to address. South Korea's enforcement record demonstrates that even the largest technology companies respond to credible legal consequences. Australia's private right of action puts enforcement power in individual hands rather than depending entirely on under-resourced regulators. The architecture of accountability is being built, piece by piece, in different jurisdictions and at different speeds, but it is being built.

The ecosystem is growing. Signal, Proton, Brave, NordVPN, Let's Encrypt, the Tor Project, and dozens of organisations less visible than these have built an infrastructure of genuine privacy protection that did not exist twenty years ago. The tools described in this book are real and they work. They are not perfect. But the gap between having them and not having them is large, and that gap is accessible to anyone reading these pages.

The culture is shifting. The language of privacy as a right rather than a preference is entering mainstream conversation in ways it was not a decade ago. People who have never heard of the Brussels Effect understand, in a way they did not before, that what happens in Europe shapes what happens globally. The generation growing up now has a more sophisticated intuition about data and consent than the generation that built the surveillance economy. That intuition, sharpened by information and experience, is a long-term force that the industry has not fully accounted for.

None of this is victory. All of it is progress. And progress, in a contest this unequal, is not a small thing.

What Remains Undetermined

The future of privacy is not settled. I want to say that clearly, because it is the most important thing I can leave you with. The surveillance economy is not invincible. It depends, as Chapter 5 described, on your passivity. It depends on you not knowing how it works. It depends on the assumption that you will not bother. Every person who bothers, in whatever way and to whatever degree they are able, is a variable the system did not fully account for.

The question of what the digital world looks like in twenty years, whether it is a world of pervasive surveillance normalised into invisibility or a world in which privacy is a genuine default rather than a premium feature, is not answered. It will be answered by the accumulation of decisions made by legislators, technologists, judges, companies, and ordinary people over the coming years.

Some of those decisions will be made by people in positions of institutional power. Many of them will be made by people who simply chose a different browser, asked a company what data it held about them, supported an organisation working on these problems, or told someone they knew what they had learned.

You are one of those people now. What you do with that is genuinely open.

The Stand Worth Making

I said at the start of this chapter that the doubt returns. I want to close with what comes with it.

The stand is worth making not because it guarantees an outcome. It does not. The stand is worth making because the alternative, deciding that the system is too large and the individual too small, and living inside that conclusion, is its own kind of surrender. And surrender, in this particular contest, does not keep you safe. It simply makes you a more compliant participant in something you did not choose.

The tools in this book reduce your exposure. The regulatory frameworks, however imperfect, establish rights you can invoke. The ecosystem of companies building differently gives you alternatives to the surveillance-funded default. And the understanding you now carry changes, in small but real ways, how you move through a world that is still being decided.

Sara left her apartment in Prague on a Tuesday morning and fourteen data points were collected before she reached the tram stop. She did not know that then. If she reads this book, she will know it now. And knowing it, she will make different choices. Not all at once. Not perfectly. But differently.

That is how this changes. One person at a time, deciding that they are not willing to be passive about it anymore. The number of people making that decision is growing. The tools available to them are better than they have ever been. The legal frameworks protecting them are more extensive than they have ever been. And the companies built to serve rather than surveil them are more numerous and more capable than they have ever been.

The future is not determined yet.

That is not a consolation. That is the point.

Thank you for reading this. Not because you chose this book, but because you stayed with it. The people who stay with difficult things are the ones who change them.

Tim Goska

Co-founder and CTO, Cyber Konekt
Prague, Czech Republic

Somewhere, a database is being updated. But somewhere else, someone just changed their default search engine. Both of these things are true. Only one of them is inevitable.



A Note From the Author

I want to tell you briefly about why this book exists, and what I built that made me want to write it.

My name is Tim Goska. I am the co-founder and CTO of Cyber Konekt, a privacy-first technology company based in Prague, Czech Republic. My co-founder Adam and I started Cyber Konekt because we kept running into the same frustration: the tools that genuinely protect your privacy online are scattered, technical, and designed for people who already know what they are doing. The people who most need them, everyday users who have not studied cybersecurity and do not want to, are the ones least likely to find and use them.

We believed, and still believe, that privacy should not require a technical degree. It should be accessible to anyone who wants it. That belief is what drove every decision in building our flagship product, and it is the same belief that drove every chapter of this book.

About VPN Toolkit

VPN Toolkit is Cyber Konekt's flagship app. It is built around the same layered defense described in Chapter 8 of this book: the idea that real privacy protection comes not from any single tool but from a sensible combination of them, working together. The app bundles four core categories of privacy protection into a single place.

VPNs, so you can browse with an encrypted connection and choose from trusted providers in one tap. Private search engines, so your queries are not logged or associated with a personal profile. Privacy-focused browsers, so tracking scripts and fingerprinting systems meet resistance rather than open access. And ad blockers, so the scripts that follow you across the web are stopped before they can run.

The reason we built it as a bundle rather than four separate tools is the same reason I structured Chapter 8 the way I did. Most people who care about privacy still do not act on it, not because they do not want to, but because the path from caring to doing involves too many decisions, too many downloads, and too much configuration. VPN Toolkit is our attempt to reduce that path to a single decision: download the app, and the layers are there.

We are also building further. A secure file Vault with encrypted storage and built-in media tools is in active development, along with a self-managed cloud backup service designed from the ground up around the principle that your files belong to you and should never pass through a server that does not.

Everything we build is guided by a principle that I hope came through in this book: that privacy is not a feature to be bolted onto a product that was designed without it. It has to be the foundation. The architecture has to assume from the first line of code that the user's data is theirs, and that our job is to protect it, not to collect it.

About Cyber Konekt


Cyber Konekt is a small team based in the Czech Republic, building a privacy-first ecosystem of tools for people who want real control over their digital lives. We fight fingerprinting, stop invasive tracking, and build products that protect your privacy across every click, search, and connection.

We are pre-revenue on several of our upcoming products and actively building toward a future where privacy-first technology is not a niche choice but the obvious one. If that mission resonates with you, we would love to have you along for it.

Cyber Konekt has affiliate relationships with certain tools and services mentioned in this book, including NordVPN, Surfshark VPN, and Incogni. These relations did not influence which tools were recommended. All products named in these pages were selected because they meet the standards described in Chapter 11, including independent auditing, transparent business models, and genuine privacy-first architecture.

VPN Toolkit: www.vpntoolkit.com

Cyber Konekt: www.cyberkonekt.com



Thank you for reading. The fact that you got this far tells me something: you are the kind of person who believes that understanding a problem is the first step toward solving it. That is the right instinct. Now go and turn the key.

Tim Goska

Co-founder and CTO, Cyber Konekt

Prague, Czech Republic

NOTES & SOURCES

Introduction:

[1] As of January 2025, 144 countries had enacted national data privacy laws, placing 82 percent of the world's population, approximately 6.64 billion people, under some form of statutory data protection.

Note: Figure from the IAPP Global Privacy Law and DPA Directory, updated January 2025. The introduction uses this as a 2026 contextual figure which is accurate as the directory had not been superseded as of the book's writing.

Source: <https://iapp.org/news/a/data-protection-and-privacy-laws-now-in-effect-in-144-countries>

[2] GDPR fines across all EU member states have exceeded 7.1 billion euros since May 2018, according to the DLA Piper GDPR Fines and Data Breach Survey published January 2026. European regulators issued approximately 1.2 billion euros in fines in 2025 alone. Ireland's Data Protection Commission accounts for 4.04 billion euros of the cumulative total.

Note: The 7.1 billion figure is the gross cumulative total across all EU jurisdictions. The Introduction and Chapter 7 cite different figures intentionally: Chapter 7 cites the Irish DPC total specifically, the Introduction cites the global EU total. Both are accurate and not contradictory.

Source: <https://www.dlapiper.com/en-us/insights/publications/2026/01/dla-piper-gdpr-fines-and-data-breach-survey-january-2026>

Chapter 1: Your Digital Shadow

[3] Four data points are enough to uniquely identify 95 percent of individuals in a mobility dataset.

Source: <https://www.media.mit.edu/projects/the-privacy-bounds-of-human-mobility/overview/>

[4] Global data broker market worth over 300 billion euros globally.

Source: <https://www.grandviewresearch.com/industry-analysis/data-broker-market-report>

[5] One major search engine processes over eight billion queries per day.

Source: <https://www.internetlivestats.com/google-search-statistics/>

[6] Shoshana Zuboff described the data economy as 'the raw material of a new economic order.'

Source: <https://www.publicaffairsbooks.com/titles/shoshana-zuboff/the-age-of-surveillance-capitalism/9781610395694/>

Chapter 2: The Business of Watching

[7] Meta reported approximately 3.56 billion daily active users across its family of apps in 2024.

Source: <https://investor.atmeta.com/investor-news/press-release-details/2026/Meta-Reports-First-Quarter-2026-Results/default.aspx>

[8] Norwegian Consumer Council 2020: ten apps transmitted data to at least 135 different third parties.

Source: <https://storage02.forbrukerradet.no/media/2020/01/mnemonic-security-test-report-v1.0.pdf>

[9] Global data broker market valued at over 300 billion euros, according to industry research.

Source: <https://www.statista.com/topics/7520/data-brokers/>

[10] Frances Haugen leak of internal Facebook documents, 2021. Algorithms amplifying divisive content. Haugen testified company was aware engagement-based system caused harm.

Source: <https://www.congress.gov/event/117th-congress/senate-event/330603/text>

Chapter 3: It Is Worse in Some Places Than Others

[11] Meta fined 1.2 billion euros by Irish Data Protection Commission in 2023 for EU-US data transfers.

Note: Largest GDPR fine issued to date at time of publication.

Source: https://www.edpb.europa.eu/news/news/2023/12-billion-euro-fine-facebook-result-edpb-binding-decision_en

[12] More than 140 countries have comprehensive privacy legislation, according to UNCTAD.

Source: <https://unctad.org/page/data-protection-and-privacy-legislation-worldwide>

[13] China's Skynet surveillance network has expanded from approximately 20 million cameras in 2018 toward a planned 570 million, incorporating facial recognition and AI-powered tracking across public spaces.

Source: <https://www.surveillancewatch.io/entities/skynet>

[14] The Sharp Eyes program, announced in 2015 with a target of 100 percent coverage of public spaces, connects surveillance cameras to residential television sets allowing neighbours to report suspicious activity to police.

Source: <https://onezero.medium.com/chinas-sharp-eyes-program-aims-to-surveil-100-of-public-space-ddc22d63e015>

[15] The Integrated Joint Operations Platform in Xinjiang, first exposed by Human Rights Watch in 2019 and upgraded to IJOP 2.0 by 2024-2025, monitors movements, communications, financial transactions and social connections of Uyghur and other Muslim minority populations. Machine learning models trained on historical detention data create feedback loops that flag the same demographic profiles repeatedly.

Source: <https://www.hrw.org/report/2019/05/01/chinas-algorithms-repression/reverse-engineering-xinjiang-police-mass-surveillance>

[16] Multiple Chinese surveillance technology companies including SenseTime, Megvii, Yitu, Dahua and Hikvision have been placed on the US Commerce Department entity list for their role in implementing surveillance against Uyghur and other Muslim minority groups in Xinjiang. Chinese-built surveillance infrastructure has been sold to governments across Africa, Central Asia and the Middle East.

Source: <https://www.atlanticcouncil.org/blogs/geotech-cues/the-west-china-and-ai-surveillance/>

Chapter 4: Real People, Real Harm

[17] Identity Theft Resource Center: 353,027,892 individuals impacted by data compromises in the US in 2023.

Note: ITRC 2023 Annual Data Breach Report.

Source: https://www.idtheftcenter.org/wp-content/uploads/2024/01/ITRC_2023-Annual-Data-Breach-Report.pdf

[18] UK Office for National Statistics: over 3.5 million incidents of fraud, year ending March 2023.

Note: ONS Crime in England and Wales bulletin.

Source:

<https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/bulletins/crimeinenglandandwales/yearendingmarch2023>

[19] New York Times 2018 investigation: location data from a data broker could track specific individuals with precision.

Note: Stuart A. Thompson and Charlie Warzel, 'Twelve Million Phones, One Dataset, Zero Privacy', New York Times, December 2018.

Source:

<https://www.nytimes.com/interactive/2019/12/19/opinion/location-tracking-cell-phone.html>

Chapter 5: Convenience Was Always the Trojan Horse

[20] Aza Raskin, widely credited with popularising infinite scroll, estimated it wastes around 200,000 human lifetimes every day.

Note: The 200,000 figure comes from Raskin's own tweet (Source). Later interviews cite a higher figure of "over a million human lifetimes a day," suggesting Raskin may have revised his estimate upward over time. The invention claim is also disputed by Microsoft engineers who hold a prior patent.

Source: <https://x.com/aza/status/1138268959982022656>

Secondary Source: <https://www.linkedin.com/pulse/full-transcript-rethink-moments-shift-your-aza-raskin-botsman>
(later interview citing the higher figure)

[21] Carnegie Mellon researchers Lorrie Faith Cranor and Aleecia McDonald: reading all privacy policies encountered in a year would consume approximately 76 work days.

Note: McDonald, A.M. and Cranor, L.F. (2008). The Cost of Reading Privacy Policies. I/S: A Journal of Law and Policy for the Information Society.

Source: <https://lorrie.cranor.org/pubs/readingPolicyCost-authorDraft.pdf>

[22] Nir Eyal, Hooked: How to Build Habit-Forming Products (2014).

Source: <https://www.nirandfar.com/hooked/>

Chapter 6: The Fingerprinting Arms Race

[23] AmIUnique study (Laperdrix et al., 2016): 81 percent of mobile browser fingerprints are unique across 118,934 browsers analysed.

Source: <https://inria.hal.science/hal-01285470/document>

[24] Fingerprinting found on more than a quarter of the top ten thousand websites.

Note: Roesner, F. et al. (2012). Detecting and Defending Against Third-Party Tracking on the Web. USENIX NSDI. Multiple subsequent studies confirm ongoing prevalence.

Source: <https://dl.acm.org/doi/10.1145/2976749.2978313>

[25] Google reversed its plan to phase out third-party cookies in July 2024.

Source: <https://privacysandbox.com/news/privacy-sandbox-update/>

[26] University of Chicago and Leibniz University Hannover, 2018: widespread misconceptions about private browsing mode. 37 percent believed employers could not track them; 22 percent believed internet providers could not see activity.

Source:

<https://dl.acm.org/doi/fullHtml/10.1145/3178876.3186088>

[27] Four websites a person visits regularly is enough to identify them with 95 percent accuracy. Published in Nature.

Note: de Montjoye, Y-A. et al. (2015). Unique in the shopping mall: On the reidentifiability of credit card metadata. Science. Related browsing study: Olejnik et al.

Source: <https://arxiv.org/pdf/2312.15489>

Chapter 7: Why the Rules Are Not Working

[28] Irish DPC 2024: 652 million euros in fines levied; 582,000 euros actually collected. Over 4 billion euros uncollected since 2018.

Source:

<https://www.irishtimes.com/business/2026/01/12/data-protection-commission-owed-more-than-4-billion-in-fines/>

[29] LinkedIn fined 310 million euros by Irish DPC in 2024.

Source: <https://www.dataprotection.ie/en/news-media/press-releases/irish-data-protection-commission-fines-linkedin-ireland-eu310-million>

[30] American Data Privacy and Protection Act passed committee 53 to 2 in July 2022.

Source: <https://www.congress.gov/crs-product/LSB10776>

Chapter 8: The Layered Defense

[31] Approximately 1.75 billion global VPN users in 2025.

Source: <https://surfshark.com/blog/vpn-users>

[32] 2024 study: 88 percent of free Android VPNs leak user data.

Source: <https://www.top10vpn.com/research/free-vpn-investigations/>

[33] It is also developed by a company whose revenues are substantially linked to advertising

Source:

https://s206.q4cdn.com/479360582/files/doc_financials/2025/q4/GOOG-10-K-2025.pdf

[34] Brave grew its user base by over 21 percent year on year.

Source: <https://brave.com/transparency/>

[35] DuckDuckGo browser passed 50 million downloads.

Source: <https://backlinko.com/duckduckgo-stats>

Chapter 9: Your Files, Your Communications

[36] IBM Cost of a Data Breach Report 2024: average cost 4.88 million USD, largest single-year increase since pandemic.

Source: <https://www.ibm.com/reports/data-breach>

[37] Snowflake breach 2024: compromised data at AT&T, Ticketmaster, and dozens of other companies.

Source: https://en.wikipedia.org/wiki/Snowflake_data_breach

Supporting Source:

<https://cloudsecurityalliance.org/blog/2025/05/07/unpacking-the-2024-snowflake-data-breach>

[38] Signal reached 70 million users in 2024.

Source: <https://www.businessofapps.com/data/signal-statistics/>

[39] WhatsApp adopted Signal Protocol in 2016.

Source: <https://signal.org/blog/whatsapp-complete/>

[40] Approximately 320 million global health app users sit outside HIPAA jurisdiction.

Note: Figure from 2024 legislative analysis of consumer health app regulatory coverage.

Source: <https://www.paubox.com/blog/how-new-privacy-rules-could-reshape-health-app-compliance>

Chapter 10: Different Threats for Different Lives

[41] 73 percent of US employers use online tracking tools to monitor websites, applications, or screen activity.

Source: <https://www.expressvpn.com/blog/workplace-surveillance-trends-us/>

[42] 59 percent of employees report stress or anxiety associated with workplace surveillance.

Source: <https://www.expressvpn.com/blog/workplace-surveillance-trends-us/>

[43] Amazon Alexa collects 28 out of 32 possible data points, 2024 Surfshark study.

Source: <https://www.globenewswire.com/news-release/2024/07/16/2913783/0/en/Study-reveals-smart-home-privacy-risks-with-Alexa-the-most-hungry-for-user-data.html>

[44] A peer-reviewed study published in 2024 by researchers at UC Davis, University College London, and Universidad Carlos III de Madrid

Source: <https://arxiv.org/html/2409.06203v1>

[45] Over 40,000 internet-connected security cameras found streaming live footage online with no password protection, accessible via a standard web browser. Some feeds found on dark web platforms selling access to live footage from homes and businesses. Note: Bitsight TRACE research, published June 2025. The 40,000 figure is described by the researchers as likely a significant underestimate.

Source: <https://www.bitsight.com/blog/bitsight-identifies-thousands-of-compromised-security-cameras>

[46] 93 percent of Americans own at least one smart home device.

Source: <https://www.ahs.com/home-matters/quick-tips/smart-home-survey/>

[47] Smart TVs use Automatic Content Recognition technology, switched on by default in every major brand as of 2026, including Samsung, LG, Sony, and Vizio.

Note: Each manufacturer uses a different name for the same technology. Samsung calls it "Viewing Information Services," LG calls it "Live Plus," Sony calls it "Samba Interactive TV," and Vizio calls it "Viewing Data."

Source: <https://www.modemguides.com/blogs/modemguides-blog/how-to-stop-smart-tv-tracking-acr-guide>

[48] ACR captures screenshots of whatever is displayed on screen, including content from external devices connected via HDMI such as laptops and gaming consoles, and transmits data to manufacturer servers. Samsung TVs transmit every minute, LG TVs every 15 seconds.

Note: Transmission rates confirmed by peer-reviewed study. On-device capture rates are alleged to be significantly more frequent in the Texas AG filings but are not confirmed by the academic study itself.

Source: <https://dl.acm.org/doi/10.1145/3646547.3689013>

[49] Study: "Watching TV with the Second-Party: A First Look at Automatic Content Recognition Tracking in Smart TVs."

Researchers at UC Davis, University College London, and Universidad Carlos III de Madrid conducted a black-box audit of network traffic leaving Samsung and LG smart TVs. Published in Proceedings of the 2024 ACM Internet Measurement Conference.

Note: Full paper available as open access PDF from UCL's research repository. UC Davis press release provides a readable summary.

Source: <https://www.ucdavis.edu/news/your-smart-tv-watching-what-you-watch> Full paper PDF: https://discovery.ucl.ac.uk/id/eprint/10200730/1/IMC_2024_ACR_Tracking_on_Smart_TV%20Camera_ready_.pdf

[50] Firmware updates on major smart TV brands silently re-enable ACR privacy settings after users have disabled them, sometimes without notification.

Note: Documented across multiple independent sources and consumer investigations. Samsung's February 2026 Texas settlement requires clearer consent prompts but does not eliminate the re-enablement issue.

Source: <https://www.smarthomeperfected.com/how-to-disable-acr-smart-tv/>

[51] Samsung settled a Texas Attorney General lawsuit over ACR data collection in February 2026, agreeing to stop collecting data without explicit consent for Texas residents and to rewrite consent screens. Lawsuits against Sony, LG, Hisense, and TCL remain active as of publication.

Note: The settlement applies only to Texas residents. Users outside Texas retain no equivalent legal protection from this specific action.

Source: <https://stateofsurveillance.org/guides/basic/disable-smart-tv-acr-surveillance-2026/>

Chapter 11: How to Spot a Privacy Tool That Is Lying to You

[52] 88 percent of free Android VPNs leak user data; 20 percent flagged as containing malware. Note: Same study as Reference 34.

Source: <https://www.top10vpn.com/research/free-vpn-investigations/>

[53] Plausible Analytics: no cookies, no personal data collection, GDPR/CCPA/PECR compliant, independent bootstrapped company.

Source: <https://plausible.io/data-policy>

[54] TelemetryDeck is a German privacy-first analytics platform founded in 2020 in Augsburg, Germany. User identifiers are double-hashed on the device before transmission, the identification date is never stored, and re-identification is mathematically impossible even with additional data. Under GDPR Recital 26, the data TelemetryDeck processes does not qualify as personal data. GDPR, CCPA, and PECR compliant. Funded by subscriptions rather than data.

Note: Integration in progress for VPN Toolkit. Chosen specifically because its architecture rules out personal data collection at a structural level rather than through policy commitments alone.

Source: <https://telemetrydeck.com/docs/guides/privacy-faq/>

Anonymisation methodology:

<https://telemetrydeck.com/docs/articles/anonymization-how-it-works/>

Chapter 12: What Good Regulation Looks Like

[55] EU AI Act came into force August 2024, world's first comprehensive AI regulation framework.

Source: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

[56] South Korea PIPC fined two US-based platforms combined 72 million USD in 2022 for behavioral data collection without consent.

Source: <https://www.lexology.com/library/detail.aspx?g=e17de49c-ba17-4e79-8773-f0085c503a54>

[57] Google was fined approximately 50 million USD by South Korea PIPC in 2022 for consent failures.

Source: https://www.theregister.com/2022/09/14/google_and_meta_fined_over/

[58] South Korea 2026 amendments: maximum penalty ceiling raised to ten percent of total global turnover; personal CEO liability introduced.

Source: <https://iapp.org/news/a/south-korea-overhauls-pipa-and-ties-fines-to-ceo-accountability>

[59] Australia Privacy Act amendments passed December 2024, introducing private right of action and raising penalties to 50 million AUD or 30 percent of turnover.

Source:

<https://www.nortonrosefulbright.com/en/knowledge/publications/be98b0ff/australian-privacy-alert-parliament-passes-major-and-meaningful-privacy-law-reform>

[60] Brussels Effect concept attributed to Anu Bradford.

Note: See Reference 13.

Source: <https://global.oup.com/academic/product/the-brussels-effect-9780190088583>

Chapter 13: The Companies Building Differently

[61] Signal Protocol adopted by WhatsApp in 2016, at which point WhatsApp had over a billion monthly active users.

Source: <https://signal.org/blog/whatsapp-complete/>

[62] Proton: founded by CERN scientists in 2014, over 100 million user accounts, 500 employees.

Source: <https://proton.me/about/team>

[64] Proton 2021: Swiss authorities ordered handover of IP address of a French climate activist involved in occupation of commercial premises.

Note: Chapter describes as 'protest-related' which is a defensible general characterisation.

Source: <https://proton.me/blog/climate-activist-arrest>

[64] Proton Foundation became primary shareholder in June 2024, with mandate to protect founding mission.

Source: <https://proton.me/blog/proton-non-profit-foundation>

[65] NordVPN: sixth independent Deloitte no-logs audit completed in late 2025.

Note: First audit conducted in 2018. ISAE 3000 framework used throughout.

Source: <https://nordvpn.com/blog/nordvpn-no-logs-assurance-engagement-2025/>

[66] Let's Encrypt: used by more than 500 million websites.

Note: Let's Encrypt's own 2025 reports indicate trajectory toward 1 billion active sites. 500 million is a conservative current estimate.

Source: <https://letsencrypt.org/stats/>

[67] Electronic Frontier Foundation: litigating for digital rights since 1990.

Source: <https://www EFF.org/about>

Chapter 14: You Are Not Powerless

Chapter 14 contains no new factual claims or statistics. All references to regulation, ecosystem companies, and tools refer back to material verified in earlier chapters. No additional references required.

GLOSSARY

This glossary defines the key terms, technologies, and regulatory frameworks used throughout the book. Where a term has multiple meanings in different contexts, the definition given here reflects how it is used in the privacy and data protection space specifically.

A

ACR (Automatic Content Recognition)

A technology embedded in most modern smart televisions that captures screenshots of whatever is displayed on screen, identifies the content using audio and visual fingerprinting, and transmits that data to the manufacturer's servers and advertising partners. ACR works regardless of whether you are watching a streaming service, a broadcast channel, or a device connected via HDMI. It operates by default on all major smart TV brands and transmits data at intervals ranging from every 15 seconds to once per minute depending on the manufacturer. Research published in 2024 by UC Davis, University College London, and Universidad Carlos III de Madrid confirmed that ACR continues to function even when the TV is used purely as an external monitor. Firmware updates can silently re-enable ACR after a user has disabled it in settings.

Ad Tech (Advertising Technology)

The infrastructure of companies, platforms, and tools that connect advertisers to audiences online. Ad tech includes the systems that track your behaviour across websites, build profiles about you, and use that information to show you targeted advertisements. The ad tech ecosystem operates largely invisibly and involves hundreds of companies

exchanging data about individual users in real time.

Algorithm

A set of rules or instructions that a computer follows to make decisions or perform tasks. In the context of this book, algorithms refer primarily to the systems used by social media platforms and search engines to decide what content to show you. These algorithms are typically optimised for engagement, which can amplify divisive or emotionally charged content because it generates more interaction.

Anonymisation

The process of removing or altering personal data so that an individual can no longer be identified from it. True anonymisation is significantly harder to achieve than most people assume. Research has consistently shown that supposedly anonymised datasets can be re-identified using a small number of additional data points. The GDPR treats anonymised data as outside its scope, but only where anonymisation is genuinely irreversible.

Artificial Intelligence (AI)

Computer systems that perform tasks typically requiring human intelligence, such as recognising images, understanding language, or making predictions. In the privacy context, AI systems are increasingly used to infer sensitive attributes from apparently innocuous data, to make consequential decisions about individuals in areas like credit scoring and employment screening, and to power the recommendation systems that shape what content people see.

Audio Fingerprinting

A tracking technique that identifies a device by the unique way its audio hardware processes sound. Like canvas fingerprinting, it works by running a small test that produces slightly different results on different devices. The result is a device identifier that persists even when cookies are deleted or private browsing is used.

B

Behavioural Data

Information about how a person acts, moves, or interacts, particularly online. This includes which websites you visit, how long you spend on each page, which links you click, how you scroll, and what you search for. Behavioural data is among the most commercially valuable categories of personal data because it is a reliable predictor of future behaviour and purchasing decisions.

Behavioural Fingerprint

A method of identifying an individual based on the patterns of their online behaviour rather than the technical characteristics of their device. Research has shown that knowing just four websites a person visits regularly is enough to identify them with 95 percent accuracy from a large anonymised dataset. Behavioural fingerprinting is harder to defeat than device fingerprinting because it is based on habits rather than hardware.

Bluetooth

A short-range wireless technology used to connect devices over distances of typically up to ten metres. In the privacy

context, Bluetooth is relevant because devices continuously broadcast a unique identifier that can be detected by sensors in shops, airports, and advertising networks to track your physical movements and presence without your awareness.

Brussels Effect

The tendency of European Union regulations to become de facto global standards because companies operating in the EU's large single market find it easier to apply EU rules everywhere rather than maintain separate compliance systems for different regions. The term was coined by political scientist Anu Bradford. The GDPR is the most significant example: many companies applied its standards globally, extending privacy rights to users in countries that had no legal requirement for them.

C

California Consumer Privacy Act (CCPA) / California Privacy Rights Act (CPRA)

The strongest privacy legislation in the United States, applying to residents of California. The CCPA, passed in 2020 and strengthened by the CPRA in 2023, gives Californians the right to know what personal data companies hold about them, the right to delete it, and the right to opt out of the sale of their data. It is broadly comparable to the GDPR in scope, though with some differences in enforcement and rights.

Canvas Fingerprinting

A tracking technique that identifies a device by the unique way its graphics hardware renders a small invisible image. Because every device renders the image slightly differently due to differences in hardware, drivers, and settings, the result functions as a unique device identifier. Canvas fingerprinting works even when cookies are deleted and in private browsing mode.

Consent

In data protection law, consent is one of the legal bases on which companies can process personal data. Under the GDPR, consent must be freely given, specific, informed, and unambiguous. A pre-ticked box, consent buried in terms and conditions, or consent that cannot be withdrawn are not valid. In practice, many consent mechanisms have been designed to look like genuine consent while falling short of this standard.

Cookie

A small file stored on your device by a website you visit. Cookies have legitimate uses, such as keeping you logged into your accounts and remembering your preferences. They are also widely used for tracking your behaviour across websites over time. Third-party tracking cookies, placed by advertising networks rather than the websites you visit directly, have been the primary mechanism for cross-site tracking and are subject to consent requirements under European law.

D

Dark Pattern

A user interface design technique that manipulates users into making choices they would not make if the options were presented fairly. In the privacy context, dark patterns are used to make it difficult to decline data collection, to bury opt-out options in multiple nested menus, or to make the 'accept all' button large and prominent while making the 'decline' option small and hard to find.

Data Broker

A company whose primary business is collecting, aggregating, and selling personal information about individuals. Data brokers typically obtain data from multiple sources, including public records, loyalty programs, social media, and purchases from other companies. They hold detailed profiles on most adults in the developed world and sell access to insurers, employers, advertisers, political campaigns, and others. Most people have never heard of the companies that hold the most detailed files about them.

Data Minimisation

A core principle of the GDPR requiring that companies collect only the personal data that is strictly necessary for the specific purpose they have stated. In practice, data minimisation is one of the principles most frequently violated by the surveillance economy, whose business model depends on collecting as much data as possible on the grounds that it might be useful in the future.

Data Portability

The right, established under the GDPR, for individuals to receive a copy of their personal data in a structured, commonly used, and machine-readable format, and to transfer it to another service provider. Data portability is designed to reduce lock-in and increase competition between services.

Data Retention

How long a company keeps personal data after it has been collected. Under the GDPR, data should not be kept for longer than is necessary for the purpose for which it was collected. In practice, many companies retain data indefinitely or for periods far exceeding what their stated purposes would justify.

DNS (Domain Name System)

The system that translates website addresses you type into the numerical IP addresses that computers use to locate each other on the internet. Every website you visit generates a DNS request, which is sent to a DNS resolver, typically operated by your internet provider by default. This means your internet provider can see every domain you visit. Using a privacy-respecting DNS resolver prevents this specific form of surveillance.

E

Encryption

The process of transforming data into a form that can only be read by someone with the correct decryption key. Encrypted data looks like meaningless noise to anyone who does not hold the key. Encryption protects data in transit and in storage, but its strength depends critically on who holds the key. Server-side encryption, where the service provider holds the key, protects against outside attackers but not against the provider itself. End-to-end encryption, where only the user holds the key, provides significantly stronger protection.

End-to-End Encryption (E2EE)

A form of encryption where messages or files are encrypted on the sender's device and can only be decrypted on the recipient's device. The service provider handling the transmission cannot read the content. Signal and Proton Mail use end-to-end encryption for messages between their users. WhatsApp also uses end-to-end encryption for message content, but collects significant metadata about who you communicate with and when.

F

Fingerprinting (Device Fingerprinting)

A method of identifying a device, and therefore its user, based on the unique combination of technical characteristics it presents to websites and services. These characteristics include screen resolution, browser version, installed fonts, timezone, language settings, and dozens of other attributes. When combined, they form a profile distinctive enough to identify most devices uniquely. Unlike cookies, fingerprints

cannot be deleted because they are calculated from the device's characteristics rather than stored on it.

G

GDPR (General Data Protection Regulation)

The European Union's comprehensive privacy regulation, which came into force on 25 May 2018. The GDPR gives EU residents the right to know what personal data is held about them, to correct it, to delete it in certain circumstances, and to object to certain types of processing. It requires companies to have a lawful basis for processing personal data, to collect only what they need, and to protect it appropriately. Fines for serious violations can reach four percent of global annual turnover or twenty million euros, whichever is higher. The GDPR is widely regarded as the global benchmark for privacy regulation, despite significant enforcement inconsistencies.

H

HIPAA (Health Insurance Portability and Accountability Act)

A US federal law that establishes privacy and security standards for protected health information held by healthcare providers, insurers, and their business associates. HIPAA is frequently misunderstood as covering all health-related data. It does not. Consumer health apps, fitness trackers, and wellness platforms are typically not covered by HIPAA because they are not healthcare providers or insurers. This creates a significant regulatory gap in which sensitive health data generated by millions of people sits outside meaningful legal protection.

HTTPS (Hypertext Transfer Protocol Secure)

The encrypted version of HTTP, the protocol used to transfer data between your browser and websites. A website using HTTPS encrypts the data exchanged between your device and the site, protecting it from interception in transit. The padlock icon in your browser's address bar indicates an HTTPS connection. HTTPS is now standard on most reputable websites, largely due to the work of Let's Encrypt, which made SSL certificates free and widely accessible.

I

IP Address (Internet Protocol Address)

A numerical label assigned to every device connected to a network. Your IP address identifies your approximate location, your internet provider, and serves as one of the primary ways websites and services identify and track you. A VPN replaces your real IP address with that of the VPN server, preventing websites from seeing your actual location or identifying you through your IP address.

L

LGPD (Lei Geral de Proteção de Dados)

Brazil's comprehensive privacy law, which came into force in 2020. The LGPD closely mirrors the GDPR in structure and principles, including requirements for lawful basis, data minimisation, and individual rights. It applies to any company processing data about people in Brazil, regardless of where the company is based.

M

Malware

Software designed to disrupt, damage, or gain unauthorised access to a computer system. In the privacy context, malware is relevant because some free applications, including a significant proportion of free VPNs, contain malicious code that can harvest data from your device, track your activity, or provide access to your device for third parties.

Metadata

Data about data. In the context of communications, metadata means not the content of what you said but who you said it to, when, how often, for how long, from where, and using what device. Metadata can be as revealing as content, sometimes more so. Knowing that you called a cancer specialist three times last week tells a story without requiring anyone to hear what was said. Most messaging apps that offer end-to-end encryption for content still collect extensive metadata.

P

PIPEDA (Personal Information Protection and Electronic Documents Act)

Canada's federal private-sector privacy law. PIPEDA requires organisations to obtain meaningful consent for the collection, use, and disclosure of personal information, and gives individuals the right to access and correct their information. Canada is currently in the process of updating PIPEDA through Bill C-27, which would significantly strengthen individual rights and enforcement powers.

PIPA (Personal Information Protection Act)

South Korea's comprehensive privacy law, administered by the Personal Information Protection Commission. Following major amendments in 2023 and 2026, PIPA is now among the most aggressively enforced privacy frameworks in the world, with fines calculated against global total revenue and personal liability provisions for senior executives introduced in 2026.

R

Real-Time Bidding (RTB)

The automated process by which advertising space on websites is auctioned to the highest bidder in the milliseconds between a user clicking a link and the page loading on their screen. During this process, data about the user, including behavioural profile, inferred interests, and demographic information, is shared with dozens or hundreds of advertisers simultaneously. Real-time bidding is the technical mechanism through which much of the surveillance economy's data exchange takes place, and it is currently under legal challenge in several jurisdictions on the grounds that it violates data protection law by design.

Right to Erasure

Also known as the right to be forgotten. Under the GDPR, individuals have the right to request that a company delete their personal data in certain circumstances, including where the data is no longer necessary for the purpose it was collected, where consent has been withdrawn, or where the data has been processed unlawfully. The right is not absolute and does not apply in all situations, but it is a genuine and

legally enforceable right that millions of people have used since 2018.

S

Signal Protocol

The encryption protocol underlying Signal's messaging system, widely regarded by cryptographers as the gold standard for secure messaging. The Signal Protocol is also used by WhatsApp for message encryption. The key difference between Signal and WhatsApp is not the encryption protocol itself but everything surrounding it: what metadata is collected, who owns the company, and what the business model is.

SSL (Secure Sockets Layer) / TLS (Transport Layer Security)

The cryptographic protocols that establish encrypted connections between your browser and websites. SSL is the older term and is now largely obsolete, having been replaced by TLS, but the two terms are often used interchangeably in everyday language. When you see a padlock icon in your browser's address bar, it indicates that the connection is protected by TLS. Let's Encrypt, a nonprofit certificate authority, has made SSL/TLS certificates free and widely available, which is largely responsible for the encrypted web becoming the standard rather than the exception. See also: HTTPS, Encryption.

Surveillance Capitalism

A term coined by Harvard professor Shoshana Zuboff to describe an economic system built on the extraction of behavioural data from human experience and its conversion into predictions about future behaviour, which are then sold to advertisers. Surveillance capitalism treats human experience as a raw material to be mined for commercial value, without the meaningful knowledge or consent of the people whose experience it is.

T

Tor (The Onion Router)

A free, open-source anonymity network that routes internet traffic through a series of volunteer-operated servers, encrypting it at each step. Tor makes it extremely difficult for anyone to trace internet activity back to the user. It is maintained by the Tor Project, a nonprofit organisation, and is used by journalists, activists, dissidents, and ordinary people who want to browse without being identified or profiled.

Tracking

The practice of monitoring an individual's online behaviour, typically across multiple websites and over time, for the purpose of building a profile that can be used for advertising or other purposes. Tracking is carried out through multiple mechanisms including cookies, device fingerprinting, IP address logging, and behavioural analysis. Most tracking happens without the user's meaningful awareness.

Two-Factor Authentication (2FA)

A security method that requires two separate forms of verification before granting access to an account: typically something you know, such as a password, and something you have, such as a code sent to your phone. Two-factor authentication significantly reduces the risk of unauthorised account access even if a password is stolen, because an attacker also needs access to the second factor.

V

VPN (Virtual Private Network)

A service that creates an encrypted tunnel between your device and a server operated by the VPN provider. From the perspective of your internet provider and the websites you visit, your traffic appears to come from the VPN server rather than your actual location. A VPN prevents your internet provider from seeing what you are doing online and hides your real IP address from the websites you visit. It does not stop fingerprinting, prevent websites from tracking your behaviour once you are on them, or protect you from apps on your phone that bypass the VPN. The business model of the VPN provider matters enormously: free VPNs frequently fund themselves by logging and selling the very data they claim to protect.

Z

Zero-Knowledge Encryption

An architecture in which a service provider stores your encrypted data but holds no key to decrypt it. This means the provider genuinely cannot access your data, even under legal compulsion. Proton Drive and Tresorit use zero-knowledge

encryption for cloud storage. The term is sometimes used loosely in marketing materials, so it is worth verifying that a product genuinely implements zero-knowledge architecture rather than simply claiming to.