A PRACTITIONER'S FIELD MANUAL

Privacy Regulations for Users of AI Models

A Practitioner's Guide for Compliance-Driven Professions

AUTHOR

Robert Bergman

FIRST EDITION • 2026

Foundations • Regulations • AI-Specific Exposures • Playbooks

Contents

Click any item to jump to that section

Privacy Regulations for Users of AI Models

A Practitioner's Guide for Compliance-Driven Professions

First Edition, 2026

Preface: How to Use This Book

This book was written for people who are accountable for compliance in their organizations and who are increasingly being asked to decide how to use artificial intelligence-based systems to improve productivity. It is also a recognition of the mess we are currently faced with in privacy regulation. Being accountable for compliance in this environment could mean choosing whether to adopt a particular AI tool. It could mean constraining employees' or external vendors' use of public models. It could unfortunately mean responding when regulated client data ends up being pasted into a chatbot by someone who should have known better. Or it could mean answering a regulator who wants to know, plainly, what the firm has done to prepare to adhere to a specific privacy regulation.

It is NOT an exhaustive legal treatise. Treatises have their place. This isn't one.

What you have in your hands is a concise, hopefully easy to use, field manual built around four things a working business owner needs in a compliance-oriented business, trying to take advantage of AI. First, a working understanding of the privacy-law landscape, including the parts of it that predate AI but govern it, anyway. Second, a simplified view of how the U.S. state-by-state patchwork creates operational friction that an AI deployment makes sharply worse. Third, the specific pressure points where AI, especially cloud-hosted or public-model AI, collides with existing privacy obligations. Fourth, a verified set of enforcement actions and lawsuits that illustrate, in concrete terms, what regulators and plaintiffs are doing when privacy rules are violated.

The book is organized so that you can read it front to back as a primer or jump to the chapter that is most relevant to your organization’s needs.

- Part I create the foundation.
- Part II walks through the regulations themselves, with the depth calibrated to how often U.S.-based practitioners encounter them.
- Part III is where the AI-specific material lives.
- Part IV contains short, action-oriented playbooks for legal, healthcare, finance, and mediation practices, followed by a chapter on standing up a minimum-viable AI privacy program.

The appendices are designed to be photocopied, pinned, or adapted. **Appendix A** (the master comparison table) is the single most useful reference in the book for most readers. If you own a highlighter, this is where to use it.

A word on scope. The book is written from a U.S. perspective. It takes international regimes seriously anyway, chiefly GDPR, the EU AI Act, and China's PIPL, which reach U.S. firms through their customers, vendors, or data flows, whether or not those firms intended them to. The 20 comprehensive U.S. state privacy laws appear here primarily as a comparative table with short profiles of the outliers. Giving each state its own full

chapter would produce a very different, much heavier book. Think of this as a sketch, not a blueprint. Please consult the table below for the abbreviations employed throughout this book. In an age that seems determined to compress entire thoughts into cryptic clusters of letters and hieroglyphics (modern emojis), the reader is gently encouraged to resist the impulse.

Part I: Foundations

Chapter 1: A Brief History of Privacy in the United States

1.1 Constitutional Roots

The United States has no general constitutional right to privacy, at least not of the kind that appears in most modern European constitutions. What it has, instead, is a Fourth Amendment, a protection against "unreasonable searches and seizures" by the government, and a set of judicially recognized privacy interests distilled over time from the Bill of Rights.

For most of the nineteenth century, Fourth Amendment doctrine was essentially about physical trespass. The government violated your privacy when its agents broke down your door, not when they listened at the wall. That model weakened as soon as technology made surveillance without physical intrusion possible. In *Olmstead v. United States* (1928), the Supreme Court allowed warrantless federal wiretaps because no trespass had occurred, over a now famous dissent from Justice Brandeis that described privacy as "the right to be let alone, the most comprehensive of rights and the right most valued by civilized men."

Forty years later, in *Katz v. United States* (1967), the Court overturned *Olmstead* and adopted what has governed since. A Fourth Amendment violation turns on whether a person has a reasonable expectation of privacy that society recognizes. *Katz* is the doctrinal point at which American privacy law stops being about personal property borderlines and starts being about information. If you have ever wondered why the Constitution keeps the government out of your email but is more relaxed about your metadata, that's the pivot.

1.2 Warren and Brandeis

The intellectual taproot of modern American privacy law is older than *Olmstead*. In 1890, Samuel Warren and the future Justice Louis Brandeis published "The Right to Privacy" in the *Harvard Law Review*. Their worry, reportedly triggered by photographic press coverage of Warren's daughter's wedding, was that a new generation of technology (cheap photography, the rotary printing press) had created harm the common law did not yet recognize. Private life, they argued, was now being exposed to mass audiences without consent.

Warren and Brandeis proposed a tort. Courts began accepting it in the early twentieth century. By mid-century, it had fractured into four distinct theories, which William Prosser crystallized in 1960.

They are:

1. **Intrusion upon seclusion.**

   Example: A journalist secretly installs a hidden camera in someone’s bedroom to record their private activities.

2. **Public disclosure of private facts**

   Example: A website publishes a person’s confidential medical diagnosis without consent, even though it’s true.

3. **False light.**

   Example: A magazine uses a person’s photo to illustrate a story about criminal activity, implying they were involved when they were not.

4. **Appropriation of name or likeness.**

   Example: A company uses a celebrity’s image in an advertisement without permission to promote its product.

These four torts are still recognized in most states. They are evidence of a century-long American tradition that treats privacy as a matter for case-by-case common-law adjudication rather than comprehensive regulation.

1.3 The Sectoral Era Begins

Comprehensive privacy legislation arrived in the U.S. first in a narrow, industry-specific structure. The Fair Credit Reporting Act, enacted in 1970, was the first major federal privacy statute of the modern era. It did not govern privacy. It governed the specific industry of consumer credit reporting, which had grown into a national system capable of destroying a person's financial life because of “errors that were easy to make and hard to correct.”

A wave of other sectoral laws followed. The Privacy Act of 1974 restricted federal agencies' handling of personal records. FERPA, passed the same year, gave parents and adult students' rights over education records. The Electronic Communications Privacy Act of 1986 extended wiretap protections into the digital age. The Video Privacy Protection Act of 1988 was famously triggered when a reporter obtained Judge Robert Bork's video-rental history during his Supreme Court confirmation fight (which is the reason some of your video-streaming history today enjoys stronger federal protection than many health and wellness apps on your phone).

Three late-1990s sectoral laws shaped the regulatory landscape AI users inherit today:

- **HIPAA (1996)** imposed privacy and security obligations on the healthcare sector. It gave us Protected Health Information (PHI) as a category and the regulatory structure under which any AI tool touching PHI must operate.

- **COPPA (1998)** regulated online collection of data from children under 13.
- **GLBA (1999)** imposed privacy and safeguarding obligations on financial institutions, giving us the concept of Nonpublic Personal Information (NPI).

Each of these is still in force. Each is still being enforced. Each is still entirely relevant to how an AI deployment is permitted, or forbidden, to handle the data in question.

1.4 The State Era and California's Leadership

For nearly two decades after GLBA, federal privacy law was essentially frozen. Congress debated comprehensive privacy bills repeatedly (the Consumer Privacy Bill of Rights in 2012, various do-not-track bills, the Consumer Online Privacy Rights Act in 2019, the American Data Privacy and Protection Act in 2022, the American Privacy Rights Act in 2024). Congress passed none of them.

Into that vacuum stepped California.

The California Constitution has, since 1972, included an express right to privacy enforceable against both state actors and private parties (Cal. Const. art. I, § 1). California passed online privacy laws throughout the 2000s and 2010s: CalOPPA in 2004, the "Shine the Light" law in 2005, and the Student Online Personal Information Protection Act in 2014. In 2018, it enacted the California Consumer Privacy Act, the first comprehensive U.S. state privacy law.

CCPA took effect on January 1, 2020. It was amended and expanded by the California Privacy Rights Act at the 2020 ballot. The CPRA became enforceable in its expanded form on July 1, 2023.

What CCPA really did was break the federal-or-nothing logjam. Once California got its act together and created a privacy regulation, other states had both a template and political cover. Virginia followed in 2021 with the VCDPA. Colorado, Utah, and Connecticut followed in 2022. By the start of 2026, twenty U.S. states had enacted comprehensive privacy laws. Every one of them has its own definitions, thresholds, rights, and penalties. None of them preempts the others. This is the landscape that we will discuss in Chapter 3.

1.5 The European Counterpoint

The European approach has differed from the start. Continental European legal systems generally treat privacy as a fundamental human right, a direct consequence of twentieth-century experience with authoritarian governments that used personal records as a weapon. The EU Data Protection Directive (1995) and its successor, the General Data Protection Regulation (2018), regulate the processing of personal data comprehensively, across every sector, with a rights-based framework and a supervisory-authority enforcement structure.

The practical consequence: a U.S. company doing business with EU residents is almost always subject to GDPR, regardless of whether it is subject to any U.S. privacy law. A multinational compliance program therefore must bridge two different regulatory philosophies. Think of it as running two thermostats in the same house. Each calibrated to a different climate.

1.6 Now enters AI

AI is not the first technology to outrun privacy law. Photography in the 1890s, wiretaps in the 1920s, databases in the 1970s, the internet in the 1990s, and social media in the early 2000s. Except for social media, each forced a reckoning of roughly the same shape.

What is distinctive about the current moment is that AI arrives in the middle of a regulatory patchwork already under strain. And it implicates every category of personal data at once: training corpora, prompt inputs, generated outputs, inference logs, and the embeddings and vector stores, and weights that sit between them.

Regulators have not written new statutes to govern most of this. What they have done is apply existing statutes (HIPAA, GLBA, FCRA, the state comprehensive laws, and GDPR) to AI as it exists. Much of the enforcement activity surveyed in Chapter 10 comprises regulators discovering, sometimes to their own surprise, that they already have the tools they need.

Chapter 2: Timeline of Privacy Regulation

The timeline below groups the regulations covered in this book into four eras. The narrative captures the through-line.

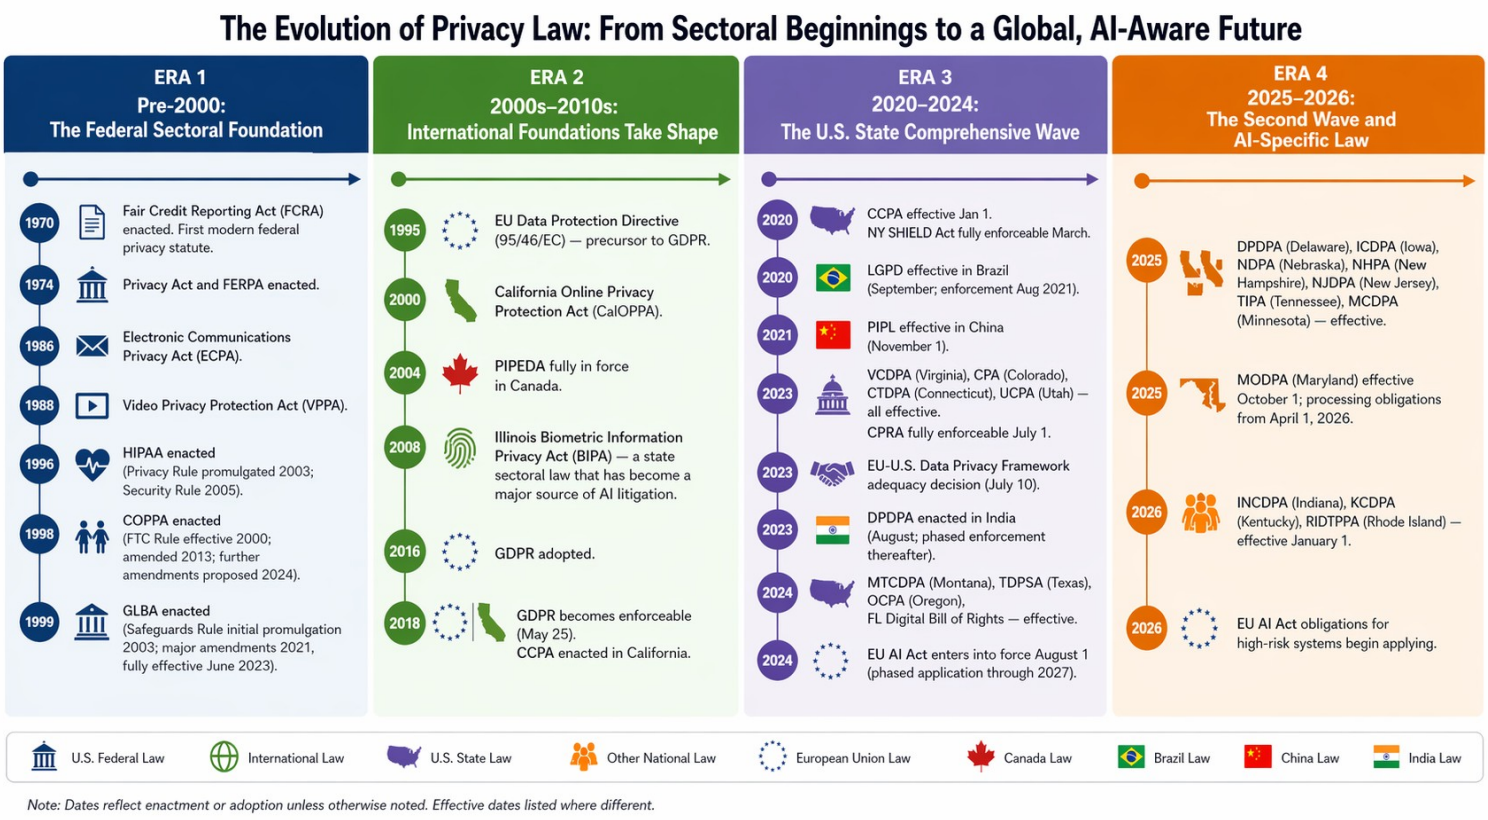

The density of the 2023 to 2026 block is not an optical illusion. In a four-year window, the United States went from one comprehensive state privacy law to twenty, and the EU added the world's first horizontal AI regulation on top of GDPR. By any reasonable measure, that is a lot of law in not very much time. Of course, if we consider that "Runaround" a science fiction short story by American writer Isaac Asimov, that featured the 3 laws of robots and ethics of machine intelligence, was written in October 1941 and first published in the March 1942, maybe it is not so dense.

Chapter 3: The Fragmentation Problem

3.1 Why the U.S. Has Twenty State Privacy Laws Instead of One Federal Law

You may finish this chapter wondering why Congress has not stepped in. The short answer: it has tried, more than once, and failed every time. In fact, there is another attempt (SECURE Data Act) traversing congress as of the date of this writing.

The most recent serious attempt was the American Privacy Rights Act, introduced in April 2024 by the chairs of the House Energy and Commerce and Senate Commerce Committees. APRA would have preempted most state comprehensive privacy laws, created a private right of action, and imposed new obligations on "covered entities," including specific rules for algorithmic decision-making. **It did not pass**.

The predecessor bill, the American Data Privacy and Protection Act, had passed out of committee in 2022 with overwhelming bipartisan support, 53 to 2, and still did not reach a floor vote. The two persistent sticking points are roughly the same each round. Scope of state-law preemption (California's congressional delegation has refused to accept a federal floor weaker than CCPA). And whether to include a private right of action.

The result is a regulatory equilibrium in which every state has concluded it must act for itself. For practitioners, this means the legal landscape you operate in depends on at least three variables: ***where your firm is located, where your clients or customers are located, and where your data is processed or stored.*** Those three variables rarely line up. When and if they do, it is usually because someone spent a great deal of money making them align.

3.2 Mapping the Differences

The twenty state comprehensive privacy laws are sometimes described as members of the "Virginia framework" because most were modeled on the VCDPA. That abbreviation is so useful, just like most abbreviations, and, naturally, it smooths over all the messy real differences just like most abbreviations. Anyway, here are the parts where they don’t match up and start to create headaches:

- **Thresholds.** Who gets covered. Most states set the trigger at 100,000 consumers whose data you process, paired with a revenue threshold. California starts higher (25 million dollars in revenue or 100,000 consumers). Texas and Nebraska apply to any business that is not a "small business" under the federal SBA definition, regardless of revenue. Delaware and Maryland set the trigger at 35,000 consumers, which is roughly an order of magnitude lower than their peers.
- **Sensitive-data scope.** Most states list sensitive data as: racial or ethnic origin, religious beliefs, mental or physical health condition or diagnosis, sex life, sexual orientation, citizenship or immigration status, genetic or biometric data, children's

data, and precise geolocation. Oregon adds transgender or nonbinary status, crime-victim status, and national origin explicitly. Maryland restricts the collection of sensitive data to what is "strictly necessary" and outright bans its sale.

- **Consumer rights.** Most states grant the classic five: access, correction, deletion, portability, and opt-out of sale or targeted advertising. Some grant an additional opt-out of profiling. Iowa does not. California and Colorado grant a broader set, including specific rights around automated decision-making. Minnesota grants a right to question a profiling decision.
- **Cure periods.** Early state laws gave businesses 30 or 60 days to "cure" a violation before facing penalties. Later laws have generally dropped cure periods or set them to sunset on a schedule. Rhode Island has no cure period at all. Montana's cure period sunsets April 1, 2026. California eliminated its cure period when CPRA took effect in 2023.
- **Universal opt-out recognition.** GPC (Global Privacy Control) is a browser-level signal that says, "do not sell or share my data." California, Colorado, Connecticut, Texas, Oregon, New Jersey, New Hampshire, Minnesota, Maryland, and Delaware all require businesses to honor GPC. Others (Virginia, Utah, Iowa, Tennessee, Rhode Island) do not.
- **Profiling and ADMT.** California's CPPA has been actively issuing regulations on automated decision-making technology since 2024. Colorado has issued parallel rules. Most other states have undefined or deferred obligations in this area, which creates unpredictability exactly where AI deployments live.
- **Penalties.** Generally, penalties are between 2,500 and 7,500 dollars per violation under Virginia-framework laws. "Per violation" sometimes means per consumer, sometimes per incident. California allows up to 7,500 per intentional or minor-involving violation. A single data-handling practice that affects a million consumers therefore has theoretical exposure in the hundreds of millions. Settled enforcement actions have generally landed in the mid-six to low-seven figures. The word "generally" is doing real work in that sentence.

3.3 The "Rights Convergence, Definitions Divergence" Problem

For AI deployments, the practical bite of state fragmentation is usually not the difference in consumer rights. It is the difference in what counts as covered data, and who counts as a covered business.

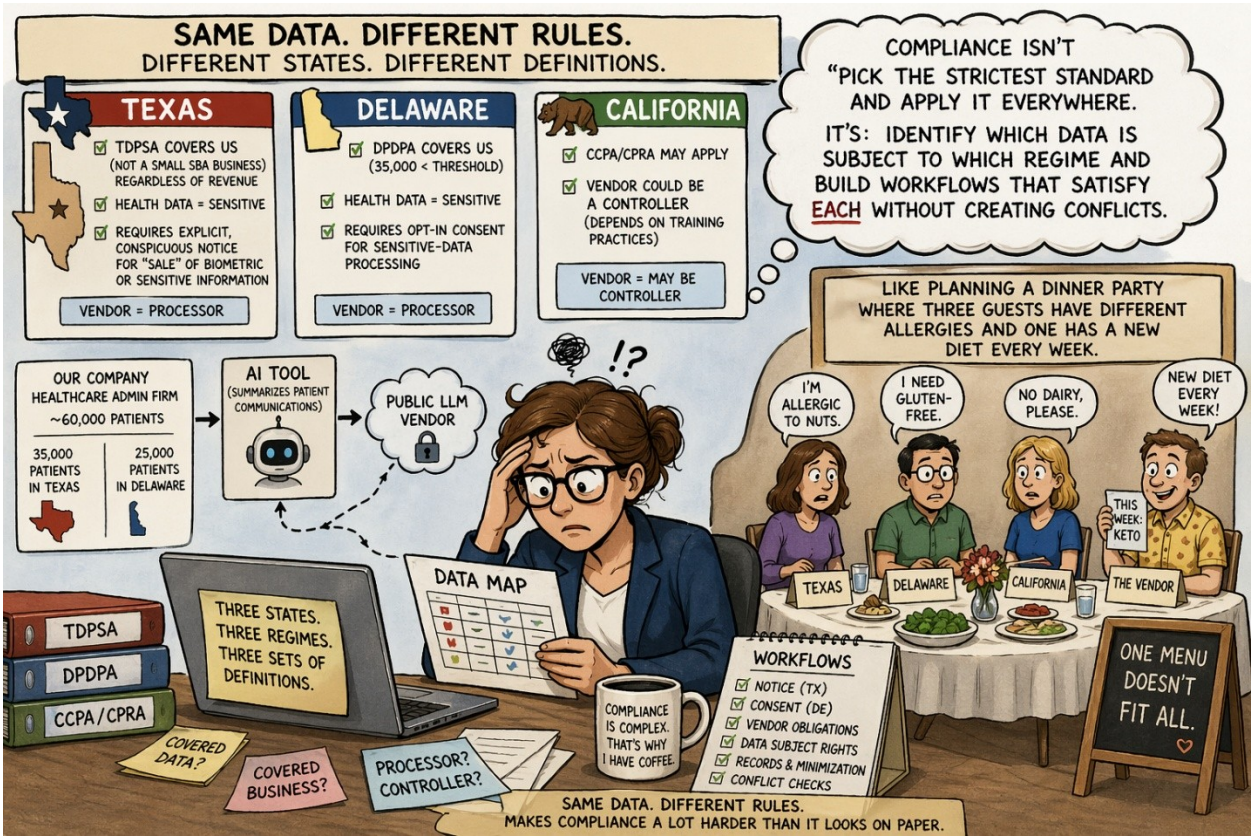

A simple example. A mid-size healthcare administration firm deploys an AI tool to help summarize patient communications. The tool routes the prompts through a public LLM vendor. The firm processes data on roughly 60,000 patients, most of whom live in two states: Texas and Delaware.

- In Texas, the firm is covered by TDPSA regardless of revenue, because it is not a small SBA business.
- In Delaware, the firm is covered by DPDPA because 35,000 is below its threshold.
- Both states consider health data sensitive. Delaware requires opt-in consent for sensitive-data processing. Texas requires explicit, conspicuous notice for the "sale" of biometric or sensitive information.
- The tool's vendor is a processor under Texas law, a processor under Delaware law, and, depending on training practices, possibly a controller under California law for the same data.

The compliance task is not "pick the strictest standard and apply it everywhere." It is "identify which data is subject to which regime and build workflows that satisfy each without creating conflicts." That is materially harder than any single statute makes it look on paper. Something like planning a dinner party where three of the guests have different allergies and one has a new diet every week.

3.4 Universal Opt-Out (GPC) Across States

Global Privacy Control is a simple idea, a header or setting in a user's browser that communicates a standing opt-out. It has become a useful proxy for a state's compliance posture.

States that require GPC recognition include California, Colorado, Connecticut, New Jersey, New Hampshire, Texas, Oregon, Minnesota, Maryland, and Delaware.

States that do not currently require it include Utah, Iowa, Tennessee, Rhode Island, Virginia, Kentucky (as enacted), and Indiana (as enacted).

California's *Sephora* enforcement (Chapter 10) was the first major action grounded in GPC non-recognition. Any business that shows a "Do Not Sell My Info" link to California users but fails to honor GPC signals carries a live exposure. You do not want to be the next one served. The suit with Sephora was settled for $1.2M.

3.5 What Fragmentation Means When You Deploy a Single AI Tool

A single AI tool, say a firm-wide deployment of a generative AI assistant, will inevitably process data about residents of multiple states. The compliance architecture most mid-size firms end up building looks like this:

1. A global acceptable-use policy governed by the strictest standard the firm faces (often GDPR or California, sometimes HIPAA if healthcare data is involved).
2. Jurisdiction-specific addenda for obligations that do not generalize, notably sensitive-data opt-in requirements, universal opt-out signal recognition, and DPIA or PIA triggers.
3. A data inventory mapping each category of data the AI tool touches to each state where the firm has consumers, and to the specific obligations that attach.
4. Vendor contracts that comply with the strictest applicable processor obligations.

Fragmentation tax is real. It is measurable. It falls disproportionately on mid-size firms that have too many consumers to ignore the state laws, but too little legal infrastructure to build per-jurisdiction workflows efficiently. The IAPP's annual benchmark has shown compliance spend rising steadily since 2020. A significant share of that rise is attributable to multi-state fragmentation specifically.

3.6 The Federal Outlook

As of April 22, 2026, with the introduction of the SECURE Act, there is finally legislation on privacy traversing congress. However, does the bill have a realistic path to enactment? No one knows for sure. The most likely near-term federal activity is probably in three narrower areas.

First, FTC enforcement under Section 5 of the FTC Act, treating deceptive or unfair privacy practices as the existing federal hook. Second, FCRA enforcement and

rulemaking applied to automated consumer-scoring tools. Third, sectoral updates, including new HIPAA guidance on AI, new GLBA Safeguards Rule enforcement, and a likely new COPPA Rule.

Part II: The Regulations in Depth

Chapter 4: Federal Sectoral Laws

4.1 HIPAA: Health Insurance Portability and Accountability Act

Who is covered. HIPAA applies to "covered entities" (health plans, healthcare clearinghouses, and healthcare providers that transmit health information electronically) and to their "business associates." A business associate is any third party that creates, receives, maintains, or transmits Protected Health Information on the covered entity's behalf. An AI vendor that processes PHI on behalf of a hospital is a business associate. **Full stop.**

What is protected. PHI, meaning individually identifiable health information, in any medium, held or transmitted by a covered entity or business associate. PHI is not just the diagnosis itself. It includes identifiers that make health information individually identifiable. The Privacy Rule's Safe Harbor de-identification standard lists eighteen specific identifiers that must be removed (see Appendix B) before health data is no longer PHI.

Core obligations. The Privacy Rule governs permissible uses and disclosures of PHI and establishes patient rights of access, amendment, and accounting of disclosures. The Security Rule requires administrative, physical, and technical safeguards for electronic PHI. The Breach Notification Rule requires notification to affected individuals, HHS, and, in large breaches, the media.

AI-specific exposures. Three recurring fact patterns account for most of the risk in practice.

- **First, PHI pasted into a public LLM**. Any time an employee of a covered entity or business associate pastes identifiable health information into a general-purpose AI tool not covered by a BAA, that is a HIPAA disclosure to an unauthorized third party. OCR has not yet published a major settlement grounded specifically on LLM prompts. The analysis, though, is identical to the Meta Pixel cases in Chapter 10. Give it time.
- **Second, AI vendors without a BAA**. HHS guidance is clear. If an AI vendor will touch PHI, you need a Business Associate Agreement. Several major cloud providers (Microsoft Azure OpenAI Service, AWS Bedrock, Google Cloud Vertex AI) will sign BAAs covering specified AI services. Consumer-facing tools (ChatGPT consumer, Gemini consumer, and the like) will not.
- **Third, training data containing PHI**. If a vendor trains a model on your PHI, that is a use beyond what is permitted under a standard BAA. Enterprise AI contracts routinely include "zero retention" and "no training on customer data" provisions. Default consumer-tier terms often do not. This is one of those places when reading the contract actually matters, rather than simply filing it.

Penalties. The HITECH Act restructured HIPAA penalties into four tiers, indexed annually for inflation. As of the August 2024 Federal Register adjustment (which carries through most of 2025 and into early 2026), the tiers (per violation; per identical-provision annual cap) are:

- **Tier 1**, no knowledge and reasonable diligence: $145 to $73,011 per violation; annual cap $ $2,190,294.
- **Tier 2**, reasonable cause, not willful neglect: $1,461 to $73,011; annual cap $ $2,190,294.
- **Tier 3**, willful neglect, corrected within 30 days: $14,602 to $73,011; annual cap $ $2,190,294.
- **Tier 4**, willful neglect, not corrected: $71,162 per violation; annual cap $2,190,294.

HHS can also impose corrective action plans lasting several years. State attorneys general have concurrent enforcement authority under HITECH. Criminal penalties under 42 U.S.C. § 1320d-6 can reach 250,000 dollars and ten years in prison for PHI obtained under false pretenses or for commercial gain. These are not theoretical; prosecutors have used them.

Practitioner takeaway. The HIPAA question is never "can we use AI?" It is "which AI, under what contract, with what data, under what controls?" A BAA is just entry level stakes. Zero-retention and no-training provisions are also table stakes. Prompt-level controls (keyword-block lists, PHI redaction at the proxy layer, or a simple prohibition on direct-to-LLM pasting) are rapidly becoming minimum requirements for any deployment that touches PHI.

4.2 GLBA: Gramm-Leach-Bliley Act

Who is covered. Financial institutions, broadly defined to include banks, credit unions, securities firms, insurance companies, financial advisors, mortgage brokers, tax preparers, auto dealers that finance vehicles, and a great many fintechs.

What is protected. Nonpublic Personal Information, meaning personally identifiable financial information provided by a consumer to a financial institution, resulting from a transaction, or otherwise obtained by the institution. NPI is a narrower category than PII. It tracks the customer relationship.

Core obligations. Two rules carry most of the practical weight.

- The *Privacy Rule* requires initial and annual privacy notices and restricts sharing of NPI with nonaffiliated third parties.
- The *Safeguards Rule* requires financial institutions to develop, implement, and maintain a comprehensive written information security program. The FTC's 2021 amendments to the Safeguards Rule, fully effective June 9, 2023, added specific

requirements: a qualified individual to oversee the program, risk assessments, access controls, encryption, multi-factor authentication, logging and monitoring, incident response plans, and board reporting. The amendments effective May 13, 2024 require notifying the FTC of any security event involving 500 or more consumers within 30 days.

AI-specific exposures. The Safeguards Rule is technology-neutral, which means AI deployments inherit all its obligations without the Rule saying "AI" anywhere.

- *Vendor due diligence*. The Rule requires financial institutions to "select service providers capable of maintaining appropriate safeguards." An AI vendor's security posture, training-data policies, and sub processor chain all become due-diligence targets.
- *Breach notification*. The 30-day FTC notification window is short. You cannot delegate that clock to a vendor's lawyers.
- *NPI in prompts*. As with HIPAA, pasting NPI into an AI tool whose contract does not cover that use is a disclosure. Depending on facts, it is also a Safeguards Rule failure.

Penalties. GLBA has no private right of action. FTC and the federal banking agencies enforce. Civil penalties under 15 U.S.C. § 45(m) for violating FTC rules (including the Safeguards Rule) can reach 53,088 dollars per violation (2025 figure, subject to annual inflation adjustment). State attorneys general have also pursued GLBA-adjacent enforcement under state consumer-protection laws.

4.3 FCRA: Fair Credit Reporting Act

Who is covered. Consumer Reporting Agencies that assemble or evaluate consumer information for third parties, the "furnishers" of information to CRAs, and "users" of consumer reports. The statute reaches beyond the traditional Big Three credit bureaus. The FTC and CFPB have both taken the position that AI-driven scoring services, tenant-screening vendors, employment background checks, and certain automated consumer-profile systems can be CRAs if their outputs are used for eligibility determinations (credit, insurance, employment, housing).

Core obligations. Accuracy (the "maximum possible accuracy" standard), permissible purpose for pulling a report, adverse action notices, dispute procedures, and reinvestigation requirements.

AI-specific exposures. This is the federal statute most likely to generate AI-driven enforcement in the near term. The reasons are structural.

- *Algorithmic scoring as a consumer report*. If an AI tool aggregates data to produce a score used for eligibility, the output may be a consumer report under FCRA. That triggers full accuracy, dispute, and notice regime.

- *Adverse action notices*. A creditor, landlord, or employer that uses an AI-generated score to decline an applicant generally must issue an adverse action notice identifying the source.
- *Tenant-screening and employment AI*. The *Louis v. SafeRent Solutions* settlement (final approval November 20, 2024; 2.275 million dollars plus injunctive terms) is the current lead case. More are in the pipeline.

Penalties. Statutory damages of 100 to 1,000 dollars per violation for willful noncompliance, plus actual damages and attorneys' fees. The CFPB can impose civil penalties up to 1,362,567 dollars per day for knowing FCRA violations (2024 figure). FCRA's private right of action has made it one of the most litigated federal statutes in the country. Some plaintiffs' firms essentially specialize in nothing else.

4.4 COPPA: Children's Online Privacy Protection Act

COPPA applies to operators of websites or online services directed to children under 13, or that knowingly collect personal information from children under 13. It requires verifiable parental consent, privacy notices, data-minimization, and restrictions on third-party sharing.

AI-specific notes. An AI tool that collects information from children triggers COPPA just as any other service would. The 2024 proposed amendments to the COPPA Rule would add new restrictions on targeted advertising and require separate consent for third-party data disclosures. Education-sector AI tools that process student data under FERPA's "school official" exception often operate under a COPPA carve-out. That carve-out has limits. It does not reach commercial uses of student data.

Penalties. Up to 53,088 dollars per violation (2025). Settlements have reached nine figures. The YouTube/Google settlement with the FTC and New York Attorney General in 2019 remains the marker at 170 million dollars.

4.5 FERPA: Family Educational Rights and Privacy Act

FERPA protects the privacy of student education records and gives parents (or students 18 and older) rights of access and consent over disclosure. It applies to educational agencies and institutions that receive federal funds.

AI-specific notes. Ed-tech AI vendors typically operate under FERPA's "school official" exception, which allows the school to share records with a contractor performing an institutional service under direct control and restricted use. An AI vendor that trains models on student data outside that carve-out creates FERPA exposure for the school.

Penalties. FERPA has no private right of action. Enforcement happens via withholding of federal funds. That remedy is theoretically dramatic and practically never invoked. The real pressure on ed-tech AI is primarily increasing from state laws (New York Education Law § 2-d, California's SOPIPA) and from institutional contract requirements.

Chapter 5: California, CCPA and CPRA

California's comprehensive privacy law is the single most consequential state privacy regime in the United States. Two reasons, really. California's economic weight, and the California Privacy Protection Agency's status as the most active state privacy regulator in the country.

5.1 Scope and Thresholds

CCPA/CPRA applies to "businesses" that (a) do business in California, (b) determine the purposes and means of processing personal information about California residents, and (c) meet at least:

- Gross annual revenue exceeding 25 million dollars.
- Annually buy, sell, or share personal information of 100,000 or more California consumers or households.
- Derive 50 percent or more annual revenue from selling or sharing personal information.

The law also defines "service providers" (processors acting on behalf of a business) and "third parties" (entities receiving personal information for their own purposes). Getting these classifications right is where most of the AI-vendor compliance work sits.

5.2 Personal Information and Sensitive Personal Information

Personal Information (PI) is defined expansively. Any information that identifies, relates to, describes, is reasonably capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular consumer or household. PI includes identifiers, commercial information, biometric information, internet activity, geolocation, sensory data, employment information, education information, and inferences drawn from any of the above.

Sensitive Personal Information (SPI), added by CPRA, is a narrower subcategory:

- Government identifiers (SSN, driver's license, passport).
- Account log-in, financial account, debit or credit card numbers with any required security code.
- Precise geolocation.
- Racial or ethnic origin, religious or philosophical beliefs, or union membership.
- Contents of mail, email, or text messages (unless the business is the intended recipient).
- Genetic data.
- Biometric information is processed to uniquely identify a consumer.

- Health information.
- Sex life or sexual orientation.

SPI unlocks an additional consumer right: the right to limit the business's use of SPI to what is reasonably necessary to perform the service.

5.3 Consumer Rights

Under CCPA/CPRA, California consumers have:

- Right to know what PI is collected, sold, or shared.
- Right to delete PI.
- Right to correct inaccurate PI.
- Right to opt out of the sale or sharing of PI (including cross-context behavioral advertising).
- Right to limit use and disclosure of SPI.
- Right to access PI in a portable format.
- Right to non-discrimination for exercising rights.
- Right to opt out of automated decision-making (scope defined by CPPA regulations, now in effect).

The CPPA has been actively rulemaking automated decision-making technology since 2024. The current regulations define ADMT broadly. They cover systems that substantially replace human decision-making in significant domains (financial services, housing, insurance, education, employment, healthcare, essential goods or services, or criminal justice). *Businesses using ADMT must provide pre-use notices, opt-out rights, and access rights to information about the logic and outcomes of decisions*.

5.4 CPPA and ADMT Regulations

The CPPA's ADMT regulations are the most aggressive state-level AI privacy rules to date. Key features:

- A "significant decision" triggers the strongest obligations.
- Pre-use notice must describe the ADMT, its intended outputs, and how consumers can exercise their rights.
- Consumers have a right to opt out of ADMT for significant decisions, subject to defined exceptions (such as when ADMT is necessary to provide a requested service and human alternatives are offered).
- Risk assessments are required for high-risk processing activities.
- Cybersecurity audits are required annually for certain businesses.

AI vendors serving California businesses will inherit most of these obligations through their service-provider contracts.

5.5 Penalties

- 2,500 dollars per violation (negligent).
- 7,500 per violation (intentional or involving a consumer under 16).
- Private right of action for breaches: 100 to 750 dollars per consumer per incident, or actual damages, whichever is greater.
- Note: CPRA eliminated CCPA's original 30-day cure period effective January 1, 2023. The AG now has discretion, not a mandatory cure window.

Violations are counted aggressively. In *Sephora*, the AG calculated penalties based on individual consumer interactions, not on a single course of conduct.

5.6 Notable Enforcement Actions

A sample here. Full treatment in Chapter 10.

- *California AG v. Sephora* (August 24, 2022), 1.2-million-dollar settlement, the first major CCPA enforcement action.
- *California AG v. DoorDash* (February 21, 2024), 375,000 dollars.
- *California AG v. Tilting Point Media* (June 18, 2024), 500,000 dollars, CCPA and COPPA, a mobile game.
- *California AG v. Healthline Media* (July 1, 2025), 1.55 million dollars, the largest CCPA settlement to date.

Chapter 6: The Twenty-State Comprehensive-Law Landscape

6.1 The Virginia Framework

Most U.S. state comprehensive privacy laws descend from the Virginia Consumer Data Protection Act. VCDPA itself drew from GDPR's controller/processor structure and from CCPA's consumer-rights list. The family traits:

- Controller and processor distinction.
- Consumer rights: access, correction, deletion, portability, opt-out of sale and targeted advertising, opt-out of profiling (usually).
- Opt-in consent for sensitive-data processing.
- Data Protection Assessments for high-risk processing.
- AG enforcement (no private right of action, in most states).
- Cure period (varying, mostly sunsetting).

Laws modeled on this framework, in rough order of enactment: VCDPA (VA), CPA (CO), UCPA (UT), CTDPA (CT), MTCDPA (MT), TDPSA (TX), OCPA (OR), TIPA (TN), FL Digital Bill of Rights, DPDPA (DE), ICDPA (IA), NDPA (NE), NHPA (NH), NJDPA (NJ), MCDPA (MN), MODPA (MD), KCDPA (KY), INCDPA (IN), RIDTPPA (RI).

Same family. Different cousins.

6.2 Comparative Snapshot

The full table lives in Appendix A. Briefly:

- **Lowest thresholds**: Delaware (35,000 consumers), Maryland (35,000), Texas and Nebraska (any non-SBA-small business).
- **No revenue threshold**: Texas and Nebraska.
- **Strictest sensitive-data regime**: Maryland (ban on sale, strict necessity for collection).
- **Broadest sensitive-data scope**: Oregon (transgender or nonbinary status, crime-victim status, and national origin all explicit).
- **No cure period**: Rhode Island; California since CPRA enforcement began; Montana's sunsets April 1, 2026; most others sunset within two to three years of effective date.
- **Mandatory GPC recognition**: California, Colorado, Connecticut, New Jersey, New Hampshire, Texas, Oregon, Minnesota, Maryland, Delaware.
- **NIST-aligned affirmative defense**: Tennessee (TIPA) is the only state that expressly offers one.

6.3 Outliers Worth Knowing

Texas (TDPSA). Applies to any person or entity that conducts business in Texas or produces products or services consumed by Texas residents, processes or sells personal data, and is not a small business under the U.S. SBA definition. The absence of a revenue threshold makes Texas the broadest state law in practical coverage. Sale of sensitive or biometric data requires explicit, conspicuous notice.

Maryland (MODPA). The most stringent U.S. state privacy law to date. Outright prohibits the sale of sensitive data. Collection of sensitive data is limited to what is "reasonably necessary and proportionate" to provide the product or service the consumer requested. No broad carve-out for loyalty programs. No broad carve-out for targeted advertising.

Oregon (OCPA). The sensitive-data definition explicitly includes transgender or nonbinary status, crime-victim status, and national origin. These categories are omitted from most peer statutes.

Tennessee (TIPA). Unique among state privacy laws in providing an affirmative defense for businesses that "create, maintain, and comply with a written privacy program that reasonably conforms" to NIST Privacy Framework or a similar program. This is the closest U.S. equivalent to ISO/IEC 27701 adoption as a compliance shield.

6.4 How Mediators and Solo Practitioners Hit Thresholds

A common misconception among small firms: "I don't have 100,000 customers, so state privacy laws don't apply to me." That assumption is wrong in several directions at once.

First, several states (Texas, Nebraska, Maryland in some contexts) do not use a 100,000-consumer threshold at all. If you are not a small SBA business (which for most professional services is under 8 to 15 million dollars in revenue and under roughly 100 employees, depending on NAICS code), you are covered in Texas and Nebraska regardless.

Second, "consumer" in these statutes is defined by residency. It counts not just customers but any individual whose personal information you process. A mediation firm in California that takes in personal information from parties across fifteen states during intake has consumers from fifteen states.

Third, many states lower the threshold if a percentage of revenue comes from selling or sharing data, or if the firm processes "sensitive data." For a mediation firm's intake records, it nearly always does.

Fourth, HIPAA, FCRA, GLBA, and state mediation-confidentiality statutes may apply regardless of the privacy-law thresholds.

Put more plainly: a solo mediator or small legal practice using an AI tool for client communications in more than two or three states should assume some state privacy

regime reaches them and should at minimum run a scoping analysis before deploying. The analysis is usually an afternoon, not a week. It is almost always shorter than the remediation.

6.5 Penalty Ranges

State statutory penalty structures, at a glance:

- Virginia: up to 7,500 dollars per violation.
- Colorado: up to 20,000 per violation (general consumer protection statute incorporation).
- Connecticut: up to 5,000 per willful violation under CUTPA.
- Texas: up to 7,500 per violation, plus injunctive relief.
- California: 2,500 / 7,500 per violation.
- Oregon: up to 7,500 per violation.
- Maryland: up to 10,000 per violation, 25,000 for repeated violations.
- Rhode Island: up to 10,000 per violation.

In practice, AG enforcement generally aggregates per-consumer or per-interaction. A data practice affecting 100,000 consumers with a 7,500 per-violation ceiling has theoretical exposure of 750 million dollars. Settled cases have generally come in much lower, in the single-digit to low-eight-figure range. The ceiling is not theoretical, though, if a state AG chooses to press.

Chapter 7: International Regulations That Affect U.S. AI Users

7.1 GDPR: General Data Protection Regulation

Why U.S. firms are subject. GDPR Article 3 gives the regulation extraterritorial reach, not extraterrestrial. It applies to any processing of personal data of individuals in the EU if the processing relates to (a) offering goods or services to those individuals (whether paid or free), or (b) monitoring behavior that takes place within the EU. A U.S. AI vendor with EU customers is subject to GDPR. A U.S. firm that uses an AI tool to process EU residents' data is subject to GDPR. A U.S. company whose marketing site is visited by a single EU resident is arguably, on some readings, also subject to GDPR, although few regulators will press that edge.

Key concepts.

- *Personal data*: any information relating to an identified or identifiable natural person. Broader than U.S. "PII." It includes online identifiers, pseudonymized data in many cases, and inference outputs.
- *Controller and processor*: the controller determines the purposes and means of processing. The processor acts on the controller's instructions. An AI vendor that trains on customer data for its own purposes becomes a controller for that data.
- *Lawful basis*: every processing activity requires a lawful basis under Article 6. The most common AI deployments are contract, legitimate interests, and consent.
- *Special categories of data* (Article 9): racial or ethnic origin, political opinions, religious beliefs, trade union membership, genetic data, biometric data, health data, sex life, sexual orientation. Processing requires a specific Article 9 condition, typically explicit consent.
- *Article 22*: rights regarding automated decisions that produce legal or similarly significant effects. Data subjects have the right not to be subject to such decisions, with exceptions for consent, contract necessity, or authorized law.

Data transfer. Since *Schrems II* (CJEU, 2020), transfers of personal data from the EU to the U.S. require an adequacy decision, Standard Contractual Clauses (SCCs) supplemented by a transfer impact assessment, Binding Corporate Rules, or another Article 46 mechanism. The EU-U.S. Data Privacy Framework (adopted July 10, 2023) provides an adequacy basis for transfers to self-certified U.S. organizations. Its durability remains uncertain. *Schrems III* litigation was filed in 2023 and is pending.

AI-specific exposures.

- *Training data*. Where is the lawful basis to train a model on European personal data? Most vendors claim legitimate interests, subject to balancing tests. EU regulators are actively examining the question. The Hamburg DPA's July 15,

2024 discussion paper on LLMs and personal data is required reading and takes the surprising position that trained models themselves may not "contain" personal data in the GDPR sense, while affirming that the training process does require a lawful basis.

- *Purpose limitation*. Article 5 requires that data collected for one purpose not be further processed in a manner incompatible with that purpose. AI training is often incompatible with the original collection purpose.
- *Data-subject rights*. Access, correction, deletion, and portability rights are difficult to satisfy against a trained model. Whether and how these rights attach to model weights themselves remains contested.
- *DPIAs*. Article 35 requires a Data Protection Impact Assessment for processing likely to result in high risk. Most AI deployments in regulated industries will trigger it.

Penalties. Two tiers under Article 83:

- Up to 10 million euros or 2 percent of global annual turnover for administrative or process violations.
- Up to 20 million euros or 4 percent of global annual turnover for core rights violations (lawful basis, data subject rights, cross-border transfer rules).

Landmark fines.

- *Meta Platforms* (Irish Data Protection Commission, May 22, 2023), 1.2 billion euros, for unlawful transfers of EU user data to the U.S. after *Schrems II*. The largest GDPR fine to date.
- *Amazon Europe Core* (Luxembourg CNPD, July 16, 2021), 746 million euros, for targeted advertising consent violations. Worth flagging: in 2026 the Luxembourg Administrative Court overturned this fine, ruling the CNPD had not adequately proved fault. The regulatory thesis is alive elsewhere in the EU; the record fine itself is not.
- *TikTok* (Irish DPC, September 15, 2023), 345 million euros, for children's data processing defaults.
- *Meta Instagram* (Irish DPC, September 2, 2022), 405 million euros, for handling of child business accounts.
- *Clearview AI* (multiple EU DPAs): France CNIL 20 million euros (October 2022), Italy Garante 20 million (March 2022), Greece HDPA 20 million (July 2022), UK ICO 7.5 million pounds (May 2022). Clearview has largely declined to pay. The DPAs have largely declined to let that stop them.

- *OpenAI* (Italian Garante): temporary ban on ChatGPT in Italy in March 2023, lifted April 2023; 15 million euro fine imposed in November 2024 for lawful-basis and transparency failures regarding training data.

7.2 EU AI Act

The EU Artificial Intelligence Act entered into force August 1, 2024. It is the first horizontal AI regulation in the world. Its provisions apply in phases through 2027.

Risk-tier structure.

- *Prohibited AI* (effective February 2, 2025): social scoring by public authorities, emotion recognition in workplaces and schools, untargeted scraping for facial-recognition databases, biometric categorization inferring sensitive attributes, and others.
- *High-risk AI* (phased; main obligations effective August 2, 2026): systems in critical infrastructure, education, employment, essential services, law enforcement, migration, and administration of justice. Obligations include risk management, data governance, technical documentation, transparency to users, human oversight, accuracy and cybersecurity, conformity assessment, and post-market monitoring.
- *Limited-risk AI*: transparency obligations (disclosure that users are interacting with AI, labeling of synthetic content).
- *Minimal-risk AI*: no obligations.

General-Purpose AI models. Specific obligations for providers of GPAI models (most frontier LLMs), including technical documentation, copyright compliance, and a summary of training data. "Systemic risk" GPAI models (currently those trained using more than 10^25 FLOPs of compute) have additional obligations, including model evaluation, systemic-risk assessment, incident reporting, and cybersecurity protection.

Penalties.

- Prohibited AI violations: up to 35 million euros or 7 percent of global turnover.
- High-risk AI violations and most other obligations: up to 15 million euros or 3 percent of global turnover.
- Providing incorrect information to authorities: up to 7.5 million euros or 1 percent of global turnover.

What this means for U.S. practitioners. The AI Act is most relevant to U.S. firms in three situations: (1) deploying an AI system whose outputs are used in the EU, (2) providing an AI system that is placed on the EU market, or (3) using an AI system whose output is used in the EU. A U.S. law firm deploying a generative AI tool firm-wide, whose EU office accesses it, may be a "deployer" of a high-risk system if the tool is

used for employment, credit, or similar decisions. That triggers obligations well beyond GDPR alone. Think of the AI Act as a new set of building codes imposed on a city that was already struggling to follow its existing zoning laws (GDPR).

7.3 EU-U.S. Data Privacy Framework

The DPF is an adequacy decision, not a law. It permits EU-to-U.S. personal data transfers to self-certified U.S. organizations without the need for SCCs, provided the recipient certifies compliance with the DPF's principles. The framework mirrors GDPR's key safeguards: purpose limitation, data minimization, security, accuracy, access, redress.

The DPF is the third U.S. adequacy decision. The first two, Safe Harbor (invalidated by *Schrems I* in 2015) and Privacy Shield (invalidated by *Schrems II* in 2020), have not aged well. A challenge to the DPF is pending. Organizations that rely on it should have SCC-based fallback documentation ready. That is not paranoia. It is maintenance.

7.4 UK GDPR

After Brexit, the UK retained GDPR as domestic law ("UK GDPR"), administered by the Information Commissioner's Office. The substance is nearly identical. The enforcement authority differs. The UK Data Protection and Digital Information Bill and successor legislation have proposed targeted reforms (removing certain DPIAs, reforming the ICO). U.S. firms with UK customers must address UK GDPR separately from EU GDPR, even if the practical controls look the same.

7.5 Canada: PIPEDA

PIPEDA is the federal Canadian privacy law. It covers private-sector handling of personal information and is built around ten fair-information principles (accountability, identifying purposes, consent, limiting collection, limiting use/disclosure/retention, accuracy, safeguards, openness, individual access, challenging compliance). Alberta, British Columbia, and Quebec have their own substantially similar laws. Quebec's Law 25 (effective in phases 2022 to 2024) is the most GDPR-like of the three.

Canada's proposed replacement statute, the Consumer Privacy Protection Act paired with an Artificial Intelligence and Data Act, has been before Parliament in various forms since 2020. It is not yet enacted as of early 2026. Watch this space.

Penalties under PIPEDA. Currently up to CAD 100,000 per violation. Under the proposed CPPA/AIDA, maximum penalties rise to CAD 25 million or 5 percent of global revenue.

7.6 Brazil: LGPD

The Lei Geral de Proteção de Dados is Brazil's general privacy law. In force since September 2020, with enforcement since August 2021. GDPR-aligned, but distinct in enforcement authority. The ANPD has become increasingly active since 2023.

Maximum administrative penalty: 2 percent of revenue in Brazil, capped at 50 million reais per infraction.

7.7 China: PIPL

The Personal Information Protection Law took effect November 1, 2021. PIPL is GDPR-inspired and, in practice, stricter on several dimensions.

- *Consent* is the primary lawful basis. Legitimate-interests equivalents are narrower.
- *Cross-border transfers* require a security assessment by the Cyberspace Administration of China, a standard contract, or certification. Thresholds trigger the CAC security assessment for large processors.
- *Sensitive personal information* has its own Article 28 regime requiring separate consent.
- *Automated decision-making* (Article 24) has its own transparency and opt-out regime.

Penalties. Up to 50 million yuan or 5 percent of prior-year turnover, plus possible suspension or license revocation. The *Didi Global* 8.026 billion yuan fine in 2022 remains the largest PIPL enforcement.

7.8 India: DPDPA

The Digital Personal Data Protection Act was enacted in August 2023. Enforcement is being phased in as implementing rules are finalized (2025 to 2026). DPDPA applies to digital personal data processed in India, and to processing outside India that relates to offering goods or services to data principals in India.

Key features:

- Consent as a primary lawful basis (with enumerated exceptions).
- Rights of access, correction, erasure, and grievance redressal.
- "Significant data fiduciary" designation for larger processors, with additional obligations.
- Cross-border transfer restrictions (rules-based).

Penalties. Up to 250 crore rupees (roughly 30 million U.S. dollars) per instance.

Part III: AI-Specific Legal Exposures

Chapter 8: Cloud-Based Servers and Privacy Law

8.1 Where the Data Lives Is Still a Legal Question

One of the more stubborn myths in compliance conversations is that cloud deployments remove geographic constraints. Nope, they do not. Most privacy regimes ask where the data is processed, where the data subject resides, and where the controller is established. Three variables. Cloud architectures sometimes decouple or hide them. They never eliminate them.

Data residency rules with current practical bite:

- *EU*: GDPR Chapter V restricts transfers of personal data to third countries without adequate protection.
- *China*: PIPL Article 38 and the CAC security-assessment regime restrict outbound transfers. Critical information infrastructure operators are generally required to store personal information within mainland China.
- *India*: DPDPA permits cross-border transfer except where restricted by government notification.
- *Russia*: Federal Law 242-FZ requires processing of Russian citizens' personal data to occur on servers physically located in Russia.
- *U.S. states*: Few hard localization rules. State laws increasingly require notice of where data is processed and sub processor disclosure.

For an AI deployment, "where is the data processed?" collapses to "where is the inference running?" In a cloud-hosted LLM, inference may be happening across several geographic regions simultaneously, depending on load-balancing. Mature AI vendors expose region-pinning controls for exactly this reason.

Of course, in the end you are using a system that doesn’t actually know anything, but just statistically guesses the next word really, well. Great, at least the vibes are geographically compliant.

8.2 The CLOUD Act Problem

The U.S. Clarifying Lawful Overseas Use of Data Act (CLOUD Act, 2018) allows U.S. law enforcement to compel U.S.-based providers to disclose data in their possession, custody, or control, regardless of where the data is physically stored. That creates a standing conflict with GDPR, which prohibits cross-border disclosure to third-country authorities except via specific Article 48 mechanisms.

Schrems II raised this conflict directly. The CJEU held that EU-to-U.S. transfers require supplementary measures to address U.S. government access risks. The DPF attempts to provide those safeguards at the adequacy-decision level. Individual contracting

parties still generally supplement SCCs with transfer impact assessments that address CLOUD Act exposure specifically.

For AI users in regulated industries, the practical question is: can my cloud AI vendor be compelled to disclose my data to U.S. law enforcement even though the data resides in the EU? For most major vendors, the answer is yes. Whether that is an acceptable risk depends on the data and the regulatory regime. Sometimes it is. Sometimes it is not. The conversation with counsel is the one you want to have *before* the subpoena arrives, not after.

8.3 Controller / Processor in an AI Context

GDPR's controller/processor framework was not designed with the current foundation LLM models in mind. Applying it is confusing at best and takes some careful planning.

- *Customer uses an AI API for a specific deployment.* The customer is the controller. The AI vendor is the processor, provided the vendor does not use inputs for its own purposes (notably training).
- *AI vendors use customer data to train or improve models.* The vendor becomes a controller for that use. This is the single most important contractual question. Is your AI vendor a controller or a processor, and what does their DPA say, not just in marketing, but in the signed document?
- *AI vendor's sub processors.* Cloud infrastructure providers, CDN providers, fine-tuning partners, safety-evaluation partners. All are sub processors under GDPR Article 28. The customer must approve them, and the vendor's DPA must pass obligations down.

HIPAA's parallel structure (covered entity, business associate, subcontractor) works similarly but with a different enforcement posture. OCR has repeatedly held covered entities responsible for business-associate failures. The 2024 Change Healthcare incident, which ultimately affected more than 100 million individuals, is the current object lesson. A single subcontractor failure can cascade through an entire healthcare payments ecosystem. Kind of like dominos, not really independent machines.

8.4 BAA-Signed AI Services

As of early 2026, the major cloud AI services that will sign HIPAA BAAs for specified workloads include:

- Microsoft Azure OpenAI Service.
- AWS Bedrock (for specified models).
- Google Cloud Vertex AI.
- Anthropic (for Claude via AWS Bedrock and via API under enterprise agreements).

- OpenAI (enterprise agreements, including ChatGPT Enterprise and API under specific configurations).

None of the free consumer tiers of these services will sign BAAs. Pasting PHI into ChatGPT consumer, Claude.ai consumer, or Gemini consumer is an unauthorized disclosure. The vendor may be excellent. The contract, though, is not the one you need.

8.5 Technical Controls Mapped to Legal Obligations

The gap between "we have a policy" and "we have a defensible compliance posture" is closed by technical controls. The ones that map most directly to legal obligations:

- ***Encryption in transit and at rest****.* Explicitly required by the Safeguards Rule. Implicitly required by most state privacy laws and by GDPR's "appropriate technical measures."
- ***Access logging and audit trails****.* HIPAA Security Rule, GLBA Safeguards Rule, SOC 2 Type 2.
- ***Key management and tenant isolation****.* Essential for multi-tenant AI services. The single largest cause of vendor risk rejection due diligence, in our experience.
- ***Prompt-level redaction or tokenization****.* A growing pattern in which PII or PHI is replaced with tokens before being sent to an LLM and re-inserted only on the client side. Reduces both legal risk and inadvertent-training risk.
- ***Confidential computing and Trusted Execution Environments****.* Increasingly relevant for cross-border workloads where the client wants cryptographic assurance that data cannot be accessed even by the cloud provider.

8.6 Three Risk Scenarios Worth Drilling On

Scenario A: U.S.-hosted AI receives EU personal data without a transfer mechanism. A U.S. AI vendor provides a SaaS product to a European customer. Data is stored and processed in U.S. regions. The vendor has not self-certified under DPF, has not executed SCCs, has not performed a transfer impact assessment. Every prompt submitted from an EU data subject is potentially an Article 46 violation. Remedies: DPF self-certification, SCCs plus TIA, or EU region hosting.

Scenario B: PHI sent to a cloud AI without a BAA. A hospital administrative team begins using a public AI tool to draft patient communications. PHI flows in prompts. No BAA exists. This is a HIPAA violation on multiple fronts (no BAA with the business associate, no accounting of disclosure, potential breach if the vendor retains data). OCR enforcement pattern: six- to seven-figure settlements plus multi-year corrective action plans.

Scenario C: NPI sent to a general-purpose LLM in violation of the Safeguards Rule. A financial advisor copies client account and portfolio details into a consumer AI tool for analysis. No DPA. No zero-retention agreement. No vendor-level safeguards.

This is a Safeguards Rule violation, and, depending on state, a state privacy-law violation. The FTC's 2024 and 2025 enforcement pipeline reflect increasing attention to exactly this fact pattern.

Three scenarios. Three different regulatory regimes. One underlying failure pattern: confusing "convenient" with "compliant."

Chapter 9: Public AI Models and Privacy Risk

9.1 Training Data Provenance

Most general-purpose LLMs released before 2024 were trained, at least in part, on data scraped from the open internet. Much of that data included personal information. Much of that personal information was scraped without the data subject's knowledge or consent. Sometimes, curated by selected engineering teams.

Under GDPR, this creates a lawful-basis problem. Vendors have generally asserted legitimate interests under Article 6(1)(f), subject to a balancing test. Since this sounds a bit confusing, here is a simple explanation: The AI companies could not realistically claim they had permission to use personal information, because they never asked. Asking three billion internet users for consent before training a model is, at best, science fiction. So, they reached for a different legal reason: a category called "legitimate interests" (formally Article 6(1)(f) of the GDPR). The way that one works is the company says, "I have a real business reason to do this, and although the people whose data I am using did not agree, my business reason outweighs the inconvenience to them." The law then requires a three-part test (purpose, necessity, balancing) to see whether that claim holds up.

Regulators are not uniformly persuaded.

The Italian Garante's action against OpenAI (March 2023 temporary ban, 15 million euro fine in November 2024) grounded its finding partly on lawful-basis and transparency failures regarding training data. The Hamburg DPA's 2024 discussion paper suggested that pre-trained models might not themselves contain personal data once training is complete, while affirming that the training process itself requires lawful basis for any personal data input. The distinction is subtle. It matters.

Under CCPA/CPRA, the "sale" and "sharing" definitions arguably capture some data flows to AI training. The CPPA's ongoing regulations on ADMT may begin to address this directly.

For a compliance-oriented business deploying AI, training-data provenance matters for two reasons. First, it creates a small but real risk of being drawn into class action as a user of a model whose training data is challenged. Second, some regulators (notably the Italian Garante and the French CNIL) have taken the position that a model's training failures make subsequent use of that model legally fragile. That is a meaningful thing to hear from the authority that will decide your next compliance audit.

9.2 Prompt and Output Risk

The second risk surface is the prompt and output stream itself.

Input risk. Prompts often contain sensitive data. Pasted contracts. Client names. Patient descriptions. Financial figures. Strategy documents. If those prompts are retained by the vendor, used for training, or accessed by vendor employees, the disclosure may violate HIPAA, GLBA, FCRA, state privacy law, or legal privilege.

Output risk. Outputs can memorize and regurgitate training data. Research on data extraction from LLMs (Carlini et al., 2021 and subsequent work) has shown that specific personal information can be extracted under certain conditions. For rare or unique data, the risk is non-negligible. For common data, it is lower. Neither risk is zero.

Inference risk. Even non-memorized outputs can reveal inferences about individuals. Medical conditions are inferred from symptoms. Credit risk inferred from demographic proxies. CCPA/CPRA expressly includes inferences drawn from personal information within the definition of PI.

9.3 Enterprise vs. Consumer Tiers

The most important distinction in AI privacy compliance is often contractual, not technical. Consumer tiers of AI tools generally include:

- Training on user inputs by default (sometimes opt-out, not always clear).
- No BAA, no DPA with teeth.
- No zero-retention guarantee.
- No tenant isolation.
- Terms of service that reserve broad rights for the vendor.

Enterprise tiers generally include:

- Zero retention, or short retention with deletion guarantees.
- No training on customer data.
- SOC 2 Type 2, ISO 27001, and often ISO 27701 or ISO/IEC 42001.
- BAA is available for specified services.
- GDPR DPA with SCCs.
- Tenant isolation.

The difference between the two is often a factor of two to five in price. It is an order of magnitude in legal exposure. For regulated industries, the enterprise tier is effectively mandatory. Treating the consumer tier as "free for work use" is a risk budget disguised as cost saving.

9.4 Zero-Retention and No-Training Options

As of early 2026:

- *OpenAI API*: no training on API customer data by default. Zero retention available on request for eligible customers.
- *Anthropic API*: no training on API customer data.
- *Google Vertex AI*: no training on customer data by default for enterprise configurations.
- *Microsoft Copilot for Microsoft 365*: prompts and responses remain within the customer's tenant, not used to train foundation models.
- *Consumer tiers* (ChatGPT Free/Plus, Claude.ai Free/Pro, Gemini): settings vary by product and jurisdiction. Default policies have shifted repeatedly. Check current terms before relying on any assumption. Really, check.

9.5 Agents, Tool Use, and Expanded Attack Surface

AI systems that can use tools, browse the web, call APIs, or operate as agents significantly expand the privacy-risk surface. Specifically:

- *Tool calls* may transmit sensitive data to third-party APIs not covered by the primary vendor's DPA.
- *Browsing* may expose client data to public webpages or create logs on web servers outside the compliance perimeter.
- *Memory* features (persistent context across sessions) create a new category of stored personal data, with retention, access, and deletion implications.
- *Multi-agent systems* may orchestrate flows across several models and vendors, each with its own data-handling posture.

The compliance analysis for an agent-based AI deployment is, in effect, the analysis for a distributed system. The same third-party risk management principles apply, just on a faster time scale.

9.6 Data-Subject Rights Against an LLM

A regulator or plaintiff asking a vendor to delete, correct, or disclose personal data held by a trained model runs into a hard technical problem. Model weights are not a database. Specific personal information is usually not locatable within weights.

For the non-technical person, model weights are the billions of internal numbers an AI system fine-tunes during training, and together they hold everything the model has learned. Think of them as the AI's memory frozen into a single file, valuable enough that companies protect them like trade secrets.

Removing memorized data requires retraining or post-hoc unlearning techniques that are not yet available.

Vendors have largely addressed this by:

- Offering deletion and correction against the prompt and output log, not the model weights.
- Asserting that weights are not personal data post-training.
- Providing opt-out mechanisms for future training.

Regulators have not uniformly accepted those positions. Several EU DPAs have ongoing investigations into precisely this question. The CPPA's ADMT rulemaking is likely to generate U.S. precedent within the next 12 to 24 months.

For practitioners, the practical advice reduces to one sentence. Don't put in what you can't accept remaining, at least in the prompt log and possibly in the model itself. Treat the prompt window the way you would treat a conversation in a hospital elevator or a courthouse hallway. You don't know who is in earshot, and the wrong listener at the wrong moment is exactly the kind of moment that ends up in a deposition.

Practical takeaways for the working professional. Most of the legal exposure in Chapter 9 comes from one bad habit: pasting client, patient, or matter material into a consumer chatbot and hoping for the best. The fix is less exotic than the risk suggests. Strip the identifiers before the text ever reaches the model, use a version of the tool that has signed a proper agreement with you, and keep a short record of what you did. The rest is hygiene.

- **Redact before you prompt.** Run the text through an application such as PIIAnomalyzer (piianomalyzer.ai), which runs locally on your computer, and replace names, dates, account numbers, medical record numbers, and addresses with placeholders. The model usually does not need real values to be useful, and a pseudonym works just as well for summarization, drafting, or analysis.
- **Use the enterprise door, not the consumer door.** ChatGPT Enterprise, Claude for Work, Gemini for Workspace, and the major API tiers offer zero-retention and no-training-on-your-data terms. The free consumer versions generally do not. For HIPAA-covered entities, that also means a signed BAA before any PHI touches the system (currently available from Anthropic, OpenAI, Google, Microsoft, and AWS Bedrock).
- **Turn off history and training in the settings.** Even on paid tiers, check that chat history, memory, and model-improvement toggles are set the way your policy requires. Defaults change, so re-check quarterly.

- **Treat prompts like discoverable work product.** Assume anything you type could end up in a deposition exhibit, a bar complaint, or an OCR investigation file. If you would not email it unencrypted, do not paste it into a chatbot.
- **Keep a light audit trail.** A one-line log (date, tool, purpose, whether PII was redacted) is enough for most compliance regimes and takes ten seconds. It is also what state bar ethics opinions (Florida, California, New York, DC) increasingly expect.
- **Match the tool to the sensitivity.** Low-risk drafting and research can go to a general enterprise chatbot. Anything involving PHI, privileged communications, trust account data, or minors' information belongs in a BAA-covered or zero-retention deployment, or a self-hosted model if the stakes are high enough.
- **Tell the client, when it matters.** ABA Formal Opinion 512 (July 2024) and most state analogues expect lawyers to consider disclosure when using generative AI on client matters. Mediators and healthcare providers face parallel expectations under their own ethics codes.

The theme across all seven bullets is the same. **Public AI is a useful tool, it is not a confidential one by default,** and the gap between those two things is where the fines and malpractice claims live. Closing the gap costs a few minutes per task and one conversation with your IT or compliance lead.

Chapter 10: Verified Legal Cases and Enforcement Actions

Every case, settlement, and fine in this chapter has been verified against primary sources (court dockets, agency press releases, or official decisions) as of publication. ***Always check current status and relevance of the case before citing.***

10.1 HIPAA and Healthcare AI

FTC v. GoodRx Holdings (February 1, 2023). First enforcement action under the FTC's Health Breach Notification Rule. GoodRx agreed to a 1.5-million-dollar civil penalty and a permanent ban on sharing health information with advertisers, after the FTC alleged GoodRx had shared user data with Meta, Google, and Criteo via tracking pixels without adequate notice. GoodRx is not itself a HIPAA-covered entity. The action nonetheless signaled the FTC's willingness to use the HBNR against digital-health services that sit just outside the HIPAA perimeter.

FTC v. BetterHelp (March 2, 2023; final approval July 12, 2023). 7.8 million dollar settlement with online therapy provider BetterHelp for sharing consumer health data (including email addresses, IP addresses, and intake-questionnaire responses indicating mental-health conditions) with Facebook, Snapchat, Criteo, and Pinterest for advertising purposes. The order bars BetterHelp from sharing consumer health data for advertising and requires deletion of transferred data. The case was the FTC's first requiring consumer refunds for health-data compromise.

FTC v. Cerebral (April 22, 2024). The structure of this settlement is more nuanced than headline summaries suggest. Cerebral agreed to a 10 million dollar civil penalty, of which only 2 million is payable due to Cerebral's inability to pay. A separate 5.1 million flows to consumer refunds, in significant part for deceptive cancellation practices. The underlying allegations: sharing sensitive data of roughly 3.2 million consumers via tracking pixels to LinkedIn, Snapchat, TikTok, and Meta. The order includes the FTC's first restriction on using health information for most advertising purposes.

Novant Health Meta Pixel Class Settlement (approved June 17, 2024). 6.66 million dollar class-action settlement (often reported as 6.6 million), in the U.S. District Court for the Middle District of North Carolina (No. 22-cv-697), approved by Judge Catherine C. Eagles. Not an MDL. It is a consolidated class action within a single district. Approximately 1.3 million patients used the portal during the class period (May 1, 2020 through August 12, 2022). Approximately 2.2 million of the settlement goes to attorneys' fees; the remainder is distributed to class members. The settlement includes a non-admission of wrongdoing provision.

HHS OCR Bulletin on Tracking Technologies (December 1, 2022; revised March 18, 2024; substantially narrowed by litigation). Not a case, but an enforcement posture. OCR's original bulletin asserted that tracking technologies on authenticated patient-portal pages typically involved PHI disclosure and required patient authorization or a BAA with the tracking vendor. In *American Hospital Association v. Becerra* (N.D.

Tex., June 20, 2024), the court invalidated the bulletin's "proscribed combination" language as exceeding HHS authority under HIPAA. HHS declined to appeal and formally withdrew its appeal on August 29, 2024. The revised March 2024 bulletin had already narrowed the guidance. IP addresses alone do not constitute individually identifiable health information, and unauthenticated webpages require actual patient health data to trigger HIPAA.

Doe v. Meta Platforms (N.D. Cal., No. 3:22-cv-03580). Consolidated class action (not an MDL) alleging Meta received health information from multiple hospital systems via the Pixel. In August 2023, Judge William Orrick III denied Meta's motion to dismiss, allowing wiretap-law and California Invasion of Privacy Act claims to proceed. Partial settlements, ongoing litigation.

10.2 GDPR and AI

Italian Garante v. OpenAI (March 2023 ban; November 2024 fine). The Italian Garante ordered a temporary ban on ChatGPT processing of Italian users' data in March 2023 over lawful-basis, age-verification, and transparency concerns. First major regulatory action against an LLM vendor in a major market. The ban was lifted in April 2023 after OpenAI implemented age-gating, transparency, and opt-out measures. On November 2024 the Garante imposed a 15 million euro fine on OpenAI for the underlying violations.

Clearview AI (2022 onward). Multiple EU data protection authorities have imposed substantial fines on Clearview AI for scraping publicly available photographs to build a facial-recognition database.

- Italy Garante, March 2022, 20 million euros.
- UK ICO, May 2022, 7.5 million pounds.
- Greece HDPA, July 2022, 20 million euros.
- France CNIL, October 2022, 20 million euros.

Clearview has contested enforcement and largely not paid. The pattern established that scraped personal data, including from public webpages, remains personal data under GDPR and requires lawful basis.

Meta 1.2 billion euro fine (Irish DPC, May 22, 2023). The largest GDPR fine to date. Imposed on Meta Platforms Ireland Limited for continuing EU-to-U.S. transfers of user data after *Schrems II*, in violation of Article 46(1). The fine and associated transfer suspension were functionally relieved by the EU-U.S. Data Privacy Framework adequacy decision the following month. The fine itself stands.

Amazon Europe Core 746 million euro fine (Luxembourg CNPD, July 16, 2021; overturned on appeal, 2026). Originally imposed for targeted advertising consent violations. The Luxembourg Administrative Court overturned the fine in 2026, ruling that

the CNPD had not adequately proved fault or negligence. Worth noting for practitioners: the underlying regulatory thesis (strict consent for targeted advertising) remains intact in parallel enforcement across other EU member states. What has changed is the benchmark fine. Not the risk posture.

TikTok 345 million euro fine (Irish DPC, September 15, 2023). For children's data processing violations, including default public settings for child accounts and the Family Pairing feature.

Meta Instagram 405 million euro fine (Irish DPC, September 2, 2022). For Instagram's handling of child business accounts, which displayed contact information publicly.

Hamburg DPA Discussion Paper on LLMs (July 15, 2024). Not an enforcement action, but an influential position paper suggesting that trained LLMs themselves may not contain personal data within GDPR's meaning (the underlying theory being that tokenization and embeddings transform text into mathematical representations that cannot be linked to specific individuals). The paper simultaneously affirms that the training process requires lawful basis and transparency. The framework has been cited in several subsequent investigations. Expect the debate to continue through at least 2026.

10.3 CCPA / CPRA Enforcement

California AG v. Sephora (August 24, 2022). 1.2 million dollar settlement. The first major CCPA enforcement action. The AG alleged Sephora failed to honor Global Privacy Control signals, did not disclose that it "sold" personal information, and did not provide compliant opt-out mechanisms. The settlement included injunctive terms requiring GPC recognition going forward. If your site shows a "Do Not Sell" link but ignores GPC, this case is your warning.

California AG v. DoorDash (February 21, 2024). 375,000 dollar settlement for CCPA violations. The AG alleged DoorDash transferred names, addresses, and transaction histories of California customers to a marketing cooperative without consumer notice, without an opportunity to opt out, and without qualifying as a service-provider relationship.

California AG v. Tilting Point Media (June 18, 2024). 500,000 dollars for CCPA and COPPA violations related to a mobile game (SpongeBob: Krusty Cook-Off). Alleged defects included a deceptive age screen that defaulted to a 1953 birth year, misconfigured SDKs that collected children's data without consent, and manipulative advertising targeted at children.

10.4 FCRA and Automated Decision-making

Louis v. SafeRent Solutions (D. Mass., final approval November 20, 2024). 2.275 million dollar class-action settlement (including up to 1.175 million in cash to class

members). The complaint alleged SafeRent's AI-driven tenant-screening algorithm disproportionately denied housing applications from Black and Hispanic applicants, and from applicants using housing vouchers, in violation of the Fair Housing Act and state consumer-protection laws. The settlement included an agreement to stop using the specific scoring algorithm for housing-voucher applicants and to make changes to scoring practices. Watch *Louis* closely. It is the first major AI-specific tenant-screening settlement under housing and FCRA theories, and more are in the pipeline.

FTC rulemaking on impersonation and negative-option AI tools (2024 onward). FTC rules and guidance documents in 2024 and 2025 increasingly treated AI-generated consumer communications under FCRA's accuracy standard where they involve eligibility determinations.

State AG enforcement on automated hiring tools. New York City Local Law 144 (2023) and adjacent state actions have produced settlements with employment-tech vendors over algorithmic bias audits and disclosure.

10.5 State AG Activity

Texas AG v. Meta biometric settlement (July 30, 2024). 1.4-billion-dollar settlement (the largest single-state AG privacy recovery ever) under Texas's Capture or Use of Biometric Identifier Act (CUBI). Distinct from TDPSA. Alleged violation: facial-geometry data collection via the "Tag Suggestions" feature without informed consent. Payment is on a five-year schedule.

Texas AG v. Google (2022, 700 million dollars). Combined multi-state settlement over location-data practices.

Connecticut AG CTDPA enforcement reports. Connecticut's AG has issued semi-annual enforcement reports since CTDPA's effective date. The early reports describe cure notices issued (over 70 in the first year) rather than imposed penalties, reflecting CTDPA's initial emphasis on voluntary correction during the cure period. That cure period has now sunset for most provisions.

Colorado AG rulemaking. Colorado has been the most active state after California in issuing rules, including rules on profiling and high-risk processing activities that have been in effect since 2023.

10.6 Training Data Class Actions

Andersen v. Stability AI Ltd. (N.D. Cal., filed January 12, 2023). Putative class action by visual artists (Sarah Andersen, Kelly McKernan, Karla Ortiz) against Stability AI, Midjourney, and DeviantArt, alleging the training of text-to-image models on copyrighted images constitutes copyright infringement and violates artists' rights of publicity. Motion-to-dismiss rulings in 2023 and 2024 narrowed but did not eliminate the claims. Case ongoing.

The New York Times Co. v. Microsoft Corp. and OpenAI, Inc. (S.D.N.Y., filed December 27, 2023). Copyright suit alleging use of millions of NYT articles in training, and that GPT-4 can regurgitate portions of NYT content verbatim. On April 4, 2025, the court denied most of defendants' motions to dismiss, allowing direct copyright infringement, contributory infringement, and trademark dilution claims to proceed. Common-law unfair-competition and "abridgment" claims were dismissed. Litigation ongoing.

Doe v. GitHub, Inc. (N.D. Cal., filed November 3, 2022). Putative class action against GitHub, Microsoft, and OpenAI over GitHub Copilot's training on open-source code without attribution. Most DMCA 1202(b) claims dismissed; breach-of-contract and open-source-license claims proceeding.

10.7 Confidentiality, Privilege, and AI

United States v. Heppner, No. 1:25-cr-00503-JSR (S.D.N.Y. February 17, 2026). Judge Jed S. Rakoff, written opinion following an oral ruling from the bench on February 10, 2026. The first federal decision in the country to hold that a criminal defendant's exchanges with a consumer-tier generative AI platform are not protected by attorney-client privilege or the work-product doctrine. For the audience of this book, it is probably the single most important case of the past several years.

Bradley Heppner, former CEO of Beneficient Company Group, was indicted in October 2025 on multiple fraud-related charges, and prosecutors seized chat records showing he had used Claude to develop defense strategies after consulting his lawyers. Heppner sought to reclaim these records under attorney-client privilege and product doctrine, but Judge Jed S. Rakoff denied the motion, ruling that privilege did not apply because Claude is not an attorney and the communications were not confidential, given the platform’s consumer-tier terms allowing data use and disclosure to third parties. The work product doctrine also failed because Heppner acted independently, without attorney direction. While the ruling is narrowly confined to unsupervised consumer AI use, it establishes a clear baseline: sharing sensitive legal information with public AI tools undermines confidentiality protections. The implications extend broadly, warning legal, healthcare, and mediation contexts that improper AI use can waive protections, while suggesting that structured, attorney-directed use under strict confidentiality agreements may still preserve them.

10.8 What These Cases Collectively Show

Three patterns stand out.

First, regulators are applying old statutes to new facts. GoodRx, BetterHelp, and Sephora, were not decided under novel AI rules. They applied the FTC Act, the HBNR, CCPA, and COPPA to new fact patterns. The lesson: you do not need AI-specific law to face AI-specific enforcement.

Second, the biggest penalties still come from the EU. Meta, TikTok, and Clearview. These dwarf any U.S. enforcement action, including the Texas biometric settlement, at the statutory level. For multinationals, GDPR remains the dominant enforcement risk by an order of magnitude.

Third, private class actions are driving the training-data law. Regulators have not written rules telling AI vendors whether training on scraped data is lawful. Courts, ruling on specific complaints brought by artists, authors, news publishers, and code authors, are writing that law one motion-to-dismiss ruling at a time.

Part IV: Practical Guidance for Compliance-Driven Professions

Chapter 11: Industry-Specific Playbooks

11.1 Legal Practice

Professional responsibility baseline. ABA Formal Opinion 512 (July 29, 2024) addresses lawyers' use of generative AI. Its core propositions: a lawyer has duties of competence, confidentiality, communication, supervision, reasonableness of fees, and candor when using AI. The ethical floor is informed client consent for disclosure of confidential information to an AI provider that does not meet strict confidentiality standards. Analogous state-bar opinions have been issued in California, Florida, New York, and others.

Privilege preservation. Pasting privileged communications into a public AI tool may waive privilege, depending on jurisdiction and facts. The cleanest analysis: work only with vendors that sign enforceable confidentiality and no-training terms; maintain a written retention policy that treats AI prompt logs as case files; avoid AI tools that retain prompts beyond what is necessary for the service. Above all, replace any PII with semantically consistent tokens or redact. When in doubt, imagine the prompt appearing on the front page of the trade press. If that bothers you, don't paste it.

E-discovery and AI. AI-assisted review is well established in e-discovery, with case law endorsing predictive coding since 2012 (*Rio Tinto v. Vale*, *Da Silva Moore*, *Moore v. Publicis Groupe*). Generative AI for document review is the current frontier. Courts have begun addressing disclosure of AI-generated work product, particularly in connection with the string of fabricated-citation sanctions that began with *Mata v. Avianca* (S.D.N.Y., 2023) and have continued through every federal circuit since.

Minimum controls. Firm-wide AUP, approved-vendor list, prohibition on pasting client information into non-approved tools, redact or replace sensitive PII, periodic training, incident-response integration.

11.2 Healthcare

HIPAA BAA checklist for AI vendors. Before deploying any AI tool that may touch PHI, confirm:

1. BAA executed, covering the specific AI service.
2. Vendor's sub processors identified and covered under flow-down provisions.
3. Zero retention, or short documented retention.
4. No training on customer data (explicit contractual prohibition).
5. Breach notification timeline consistent with 45 C.F.R. § 164.410 (no more than 60 days; typically 5 to 15 days contractually).
6. Audit and termination rights.

7. Geographic processing limits, if applicable.
8. Vendor's SOC 2 Type 2, HITRUST, or comparable third-party attestation.

PHI de-identification in prompts. Where a use case does not require identifiers, implement redaction or replacement tokenization at the prompt layer. Several healthcare AI middleware vendors offer this as a turnkey service.

State health-data laws. Washington's My Health My Data Act (MHMDA, effective March 31, 2024) extends consumer-health-data protection beyond HIPAA, covering many digital-health companies and tools outside the traditional HIPAA perimeter. Nevada has enacted parallel legislation. Other states are considering similar bills.

State medical privacy statutes. California's CMIA, New York's SHIELD Act (in its medical context), Texas's Medical Records Privacy Act, and other state-level statutes impose additional requirements. Many carry private rights of action that HIPAA lacks.

11.3 Finance

Safeguards Rule compliance for AI deployments. The FTC's amended Safeguards Rule is technology-neutral. An AI deployment inherits risk-assessment, access-control, encryption, logging, incident-response, and MFA obligations. The qualified individual overseeing the information-security program needs to own the AI risk assessment.

SEC and FINRA guidance. The SEC's 2024 proposed rule on broker-dealer and investment-adviser conflicts of interest arising from predictive data analytics (including AI) has not been finalized. It signals a direction, though. FINRA has issued generative AI notices addressing supervision, recordkeeping, and communications with the public.

FCRA exposure. AI models used for credit underwriting, scoring, or eligibility determinations fall squarely within FCRA. Accuracy, adverse-action notice, and dispute-handling obligations apply to the AI output.

Anti-money-laundering and AI. BSA/AML programs increasingly incorporate AI for transaction monitoring. FinCEN guidance emphasizes model-governance obligations analogous to OCC Model Risk Management Guidance (SR 11-7).

State insurance departments. The NAIC's Model Bulletin on the Use of Artificial Intelligence Systems by Insurers (2023) has been adopted in a majority of states. It requires an AI System Program and governance analogous to SR 11-7.

11.4 Mediation, ADR and ODR

Mediation practice is one of the least-covered compliance areas in the AI discussion. The risks are sharp.

Confidentiality is the foundation. The Uniform Mediation Act (adopted in twelve states) and state-specific mediation statutes establish strong confidentiality protections for mediation communications. Those protections are generally not waived by sharing

within a mediation. **They are often waived by sharing with an AI vendor that does not qualify as a mediation participant**. While other standards have been proposed for ODR (The National Center for Technology and Dispute Resolution (NCTDR) and the International Council for Online Dispute Resolution (ICODR), ISO), they are only standards proposals and not enforceable law.

Public LLMs and mediation communications. A mediator who pastes party statements into a consumer AI tool to draft a memorandum has plausibly waived mediation confidentiality as to the prompted content. There is not yet reported case law directly on this. The analysis, though, follows straightforwardly from existing mediation-confidentiality doctrine.

Recommended posture. Mediators should use only enterprise-tier AI services with documented no-training and no-retention terms, or relatively closed systems like Next Level Mediation. Incorporate AI use into mediation agreements and informed-consent forms. Prefer on-device or self-hosted tools where feasible and tokenize any PII information before uploading document / prompts to AI. Treat AI prompt logs as mediation records for retention-policy purposes.

Chapter 12: Building an AI Privacy Program

12.1 The Minimum Compliance Stack

A firm that wants to adopt AI responsibly needs, at minimum, seven elements.

1. **Data inventory.** What personal information does the firm hold, and which categories are implicated by which AI deployments?
2. **Data protection impact assessment (DPIA/PIA).** Required under GDPR for high-risk processing. Advisable under CCPA/CPRA. Increasingly required under state laws. DPIA becomes the evidentiary anchor for your risk-based approach.
3. **Vendor due diligence.** Security posture, DPA, training policies, sub processors, audit rights, geographic controls. (See Appendix D for a sample questionnaire.)
4. **Acceptable-use policy.** Employee-facing policy describing approved tools, prohibited uses, data-handling rules, and consequences for noncompliance. (See Appendix C.)
5. **Employee training.** Annual at minimum. Incident-driven updates as enforcement patterns change.
6. **Monitoring and logging.** Technical controls to detect policy violations, including outbound-data monitoring on approved and unapproved AI tools.
7. **Incident response.** A plan that has actually been tabletop-exercised. One that treats AI-related data incidents as potential breaches under applicable statutes, with 30- to 60-day regulatory notification timelines built in.

12.2 NIST AI RMF and ISO/IEC 42001

For firms seeking a recognized program framework:

- **NIST AI RMF** (1.0, January 2023). Organized around Govern, Map, Measure, and Manage functions. Free, voluntary, widely adopted in the U.S. Tennessee's TIPA provides an affirmative defense for programs that align with NIST frameworks.
- **ISO/IEC 42001** (2023). AI Management System Standard. Certifiable through third-party audit. Adopts a structure familiar from ISO 27001.

For small and mid-size firms, NIST AI RMF is usually the right starting point. Larger firms, and firms operating in EU markets, often move toward ISO/IEC 42001 certification.

12.3 One-Page Employee Checklist

A practical artifact. Pin it wherever employees use AI.

- Is the tool on the firm's approved-vendor list? (If no, stop.)

- Does the content include any of the following: client names, patient information, account numbers, proprietary matter details, privileged communications, or mediation communications? (If yes, redact or stop.)
- Does the output need to be independently verified before being relied on or sent to a client? (Almost always yes.)
- Will you retain the prompt and output as part of the client or matter record? (Usually yes.)
- If the output informs a decision with legal effect (employment, credit, housing, healthcare), has the decision been reviewed by a human authorized to make it? (Required under most applicable regimes.)

Appendix A: Master Comparison Table

The table below summarizes all 34 regulations covered in this book. In the printed edition it appears as a fold-out reference.

#	Regulation	Jurisdiction	Effective	Core Subject Matter	Max Penalty	AI-Relevant Flags
1	**COPPA**	U.S. Federal	1998 (2013 amends)	Online data from under 13	$53,088 per violation	Consumer-facing AI with minors
2	**HIPAA**	U.S. Federal	1996 (Privacy Rule 2003)	PHI	Tier 4: $71,162 per violation, $2.13M annual cap (2024)	BAA required for AI vendors
3	**GLBA**	U.S. Federal	1999 (Safeguards 2023)	NPI	~$53,088 per violation	Safeguards Rule applies to vendor AI
4	**FCRA**	U.S. Federal	1970	Consumer reports	$100 to $1,000 statutory; CFPB daily penalties	Algorithmic scoring
5	**FERPA**	U.S. Federal	1974	Education records	Funding withholding	Ed-tech AI
6	**NY SHIELD**	NY	2020	Breach plus security	$5,000 per violation	Reasonable-security baseline
7	**CCPA / CPRA**	CA	2020 / 2023	Broad consumer privacy	$2,500 / $7,500 per violation	ADMT rules
8	**VCDPA**	VA	2023	Personal / sensitive data	$7,500 per violation	Profiling opt-out
9	**CPA**	CO	2023	Personal / sensitive data	$20,000 per violation	High-risk processing rules
10	UCPA	UT	2023	Narrower scope	$7,500 per violation	Limited profiling rights
11	CTDPA	CT	2023	Virginia framework plus kids	$5,000 per willful violation	GPC required

#	Regulation	Jurisdiction	Effective	Core Subject Matter	Max Penalty	AI-Relevant Flags
12	MTCDPA	MT	2024	Virginia framework	$10,000 per violation	Cure sunsets April 2026
13	TIPA	TN	2025	Virginia framework	$7,500 per violation	NIST affirmative defense
14	OCPA	OR	2024	Broad sensitive data	$7,500 per violation	Broadest sensitive scope
15	TDPSA	TX	2024	No revenue threshold	$7,500 per violation	Any business in state
16	ICDPA	IA	2025	Narrower rights	$7,500 per violation	No profiling opt-out
17	INCDPA	IN	2026	Virginia framework plus DPIA	$7,500 per violation	DPIA for heightened risk
18	DPDPA	DE	2025	Low threshold (35K)	$10,000 per violation	Expanded teen protections
19	NDPA	NE	2025	Texas-like, non-SBA	$7,500 per violation	Broad coverage
20	NHPA	NH	2025	Virginia framework	$10,000 per violation	GPC required
21	NJDPA	NJ	2025	Broad "sale" definition	$10,000 first / $20,000 subsequent	No revenue percent threshold
22	KCDPA	KY	2026	Virginia framework	$7,500 per violation	Close to VCDPA
23	MCDPA	MN	2025	Adds profiling question	$7,500 per violation	Right to question profiling
24	MODPA	MD	2025 (2026 obligations)	Strictest U.S. state	$10,000 / $25,000 repeat	Bans sale of sensitive data
25	RIDTPPA	RI	2026	No cure, no GPC	$10,000 per	Disclosure-focused

#	Regulation	Jurisdiction	Effective	Core Subject Matter	Max Penalty	AI-Relevant Flags
				requirement	violation	
26	GDPR	EU	2018	All personal data	€20M or 4% global	Article 22 ADM
27	DSA	EU	2022 / 2024	Platform content	Up to 6% global turnover	Targeted advertising limits
28	DMA	EU	2022 / 2023	Gatekeeper conduct	Up to 10% global turnover	Cross-context data limits
29	EU-US DPF	EU-US	2023	Transfer mechanism	Varies	Adequacy, not a law
30	EU AI Act	EU	2024 (phased)	AI system risk tiers	€35M or 7% global	Prohibited plus high-risk tiers
31	LGPD	Brazil	2020 / 2021	GDPR-aligned	2% of BR revenue, R$50M cap	Sensitive-data regime
32	PIPEDA	Canada	2004	Private-sector	CAD $100K (current)	CPPA/AIDA pending
33	PIPL	China	2021	Broad plus transfer controls	RMB 50M or 5%	Strict consent, sensitive-PI regime
34	DPDPA	India	2023 (phased)	Digital personal data	₹250 crore per instance	Children's-data rules

Appendix B: HIPAA Safe Harbor Identifiers

The eighteen identifiers that must be removed for PHI to be de-identified under the Safe Harbor method (45 C.F.R. § 164.514(b)(2)):

1. Names.
2. All geographic subdivisions smaller than a state, including street address, city, county, precinct, ZIP code (and equivalent geocodes), except initial three digits of ZIP under specific rules.
3. All elements of dates (except year) directly related to an individual, including birth date, admission date, discharge date, death date; all ages over 89 and all elements of dates indicative of such age (aggregate into 90 and older).
4. Telephone numbers.
5. Vehicle identifiers and serial numbers, including license plate numbers.
6. Fax numbers.
7. Device identifiers and serial numbers.
8. Email addresses.
9. Web URLs.
10. Social Security numbers.
11. IP addresses.
12. Medical record numbers.
13. Biometric identifiers, including finger and voice prints.
14. Health plan beneficiary numbers.
15. Full-face photographs and any comparable images.
16. Account numbers.
17. Any other unique identifying number, characteristic, or code.
18. Certificate or license numbers.

Plus: the covered entity must have no actual knowledge that the remaining information could be used to identify an individual.

Appendix C: Example Acceptable Use Policy for AI Tools

Adapt to firm size, regulatory posture, and local counsel review.

[FIRM NAME]. Artificial Intelligence Acceptable Use Policy.

1. Purpose and Scope. This policy governs use of artificial intelligence tools ("AI Tools") by all personnel of [FIRM NAME]. It applies to all AI Tools accessed in connection with firm business, whether licensed by the firm, licensed by an employee, or publicly available.

2. Approved Tools. Only AI Tools on the firm's Approved Vendor List may be used for work that involves firm, client, patient, or customer information. The Approved Vendor List is maintained by the [designated role] and is available at [location]. Use of any AI Tool not on the list requires written approval from the [designated role].

3. Prohibited Uses. Personnel shall not enter any of the following into an AI Tool, regardless of whether the tool is on the Approved Vendor List, unless expressly authorized in writing for the specific use:

- Client, patient, or customer identifying information.
- Protected Health Information (PHI).
- Nonpublic Personal Information (NPI) under GLBA.
- Consumer report information under FCRA.
- Personal data of EU, UK, or other non-U.S. individuals subject to extraterritorial privacy regimes.
- Mediation communications and privileged attorney communications.
- Firm-proprietary or confidential information.
- Trade secrets, regardless of source.

4. Required Workflow.

a. Before entering any regulated data into an approved AI Tool, redact identifiers or use firm-provided tokenization where available.

b. Prompt and output logs are firm records. They shall be retained and disposed of per the firm's records retention schedule.

c. AI-generated outputs used in client, patient, or customer deliverables shall be independently reviewed by a qualified professional before release.

d. Decisions with legal effect on individuals (employment, credit, housing, healthcare, legal representation) shall not be made solely by an AI Tool. Human review is required.

5. Vendor Due Diligence. Adding an AI Tool to the Approved Vendor List requires completion of the firm's vendor privacy and security questionnaire; review of the vendor's DPA, BAA (where applicable), and sub processor list; and approval by [designated role].

6. Training. All personnel shall complete AI privacy and security training on onboarding and annually thereafter.

7. Incident Reporting. Any suspected violation of this policy, or any inadvertent disclosure of regulated data to an AI Tool, shall be reported to the [designated role] immediately. Incidents shall be assessed for breach-notification obligations under applicable law.

8. Enforcement. Violations may result in discipline up to and including termination. They may be reported to licensing authorities or law enforcement where required.

9. Review. This policy shall be reviewed at least annually and updated as regulatory guidance, enforcement patterns, and firm deployments evolve.

Approved by: _____ Date: _____

Appendix D: Vendor Due Diligence Questionnaire Framework

Company and Service

1. Legal entity name, headquarters, and jurisdictions of incorporation.
2. Specific AI service(s) proposed for our use.
3. Model(s) underlying the service, including provider if different from vendor.

Data Handling

4. Will our data be used to train or fine-tune any model? If yes, on what basis and with what opt-out?
5. Default retention period for prompts, outputs, logs, and embeddings. Ability to configure shorter retention.
6. Data residency options. Can processing be pinned to specified regions?
7. Sub processor list. Change-notification process.

Legal and Compliance

8. Standard Data Processing Agreement terms (attach).
9. Willingness to execute a Business Associate Agreement for HIPAA-covered use (attach template).
10. SCCs or DPF self-certification for cross-border transfers.
11. Representations regarding lawful basis for training data.

Security

12. Third-party attestations (SOC 2 Type 2, ISO 27001, HITRUST, ISO/IEC 42001).
13. Encryption in transit and at rest. Key management approach.
14. Tenant isolation architecture.
15. Access controls for vendor employees.
16. Incident response. Notification SLA. Historical incident disclosure.

Technical AI Controls

17. Prompt-injection protections.
18. Safeguards against extraction of memorized training data.
19. Audit logging granularity.
20. Customer controls for model version pinning.

Business

21. Financial stability, insurance coverage, audit rights, termination, data-return provisions.

Appendix E: Glossary

Business Associate. Under HIPAA, a person or entity that performs functions involving PHI on behalf of a covered entity, subject to a Business Associate Agreement (BAA).

Controller. Under GDPR, the entity that determines the purposes and means of processing personal data.

Data Protection Impact Assessment (DPIA). Required under GDPR Article 35 for high-risk processing.

Deployer. Under EU AI Act, a party using an AI system under its authority (distinguished from the provider).

General-Purpose AI Model (GPAI). Under EU AI Act, a model trained on a large amount of data using self-supervision at scale, displaying significant generality and capable of competently performing a wide range of distinct tasks.

Global Privacy Control (GPC). A browser-level signal indicating a user's general opt-out preference.

Nonpublic Personal Information (NPI). GLBA's defined category of protected financial information.

Processor. Under GDPR, the entity that processes personal data on behalf of a controller.

Protected Health Information (PHI). HIPAA's defined category of individually identifiable health information.

Service Provider. Under CCPA/CPRA, a processor acting on behalf of a business, subject to specific contractual restrictions.

Sensitive Personal Information (SPI). CPRA-added subcategory of PI with additional protections.

Appendix F: Source List and Further Reading

Primary sources.

- HIPAA Privacy, Security, and Breach Notification Rules. 45 C.F.R. parts 160 and 164.
- HIPAA Civil Monetary Penalty inflation adjustment. 89 Fed. Reg. 64818 (August 8, 2024).
- GLBA Safeguards Rule. 16 C.F.R. part 314.
- GLBA Safeguards Rule amendments. 86 Fed. Reg. 70272 (December 9, 2021) and 89 Fed. Reg. 47072 (May 13, 2024).
- FCRA. 15 U.S.C. § 1681 et seq.
- CCPA / CPRA. Cal. Civ. Code § 1798.100 et seq.; CPPA regulations at 11 C.C.R. §§ 7000 through 7400.
- GDPR. Regulation (EU) 2016/679.
- EU AI Act. Regulation (EU) 2024/1689.
- ABA Formal Opinion 512 (July 29, 2024).

Secondary sources and guidance.

- HHS OCR Bulletin on Tracking Technologies (December 1, 2022; revised March 18, 2024), as narrowed by *American Hospital Association v. Becerra* (N.D. Tex., June 20, 2024).
- Hamburg DPA, Discussion Paper, "Large Language Models and Personal Data" (July 15, 2024).
- Italian Garante, Provision against OpenAI (November 2024).
- CPPA ADMT Regulations (2024 to 2025 rulemaking).
- NIST AI RMF 1.0 (January 2023).
- ISO/IEC 42001:2023.

Enforcement action sources.

- FTC press releases and consent orders (ftc.gov/news-events).
- California Attorney General settlement announcements (oag.ca.gov).
- Irish Data Protection Commission decisions (dataprotection.ie).
- European Data Protection Board coordinated decisions (edpb.europa.eu).
- IAPP Global Privacy Directory (iapp.org).