THE SUBTLE ART OF GIVING A F**** ABOUT DATA PRIVACY

Everything Developers Need to Know About Privacy Compliance

TONIC
THE FAKE DATA COMPANY

TABLE OF CONTENTS

Aspects of data privacy covered in this ebook

This ebook identifies major US and international privacy laws, summarizes common themes, presents penalties, and suggests best practices in avoiding compliance violations for software developers.

We provide an overview of the essentials so you won't get bogged down in legalese. Links to the full laws are included as a reference tool.

This information is based on a data privacy webinar sponsored by Tonic, conducted by Justin P. Webb, Data Privacy and Cybersecurity Attorney, CISO, at the law firm Godfrey & Kahn.

Although privacy laws are constantly in flux, clear, actionable steps exist to protect personal data from the risk of exposure and achieve compliance across the board, regardless of where your data lives or where it needs to go to reach your developers.

Let's dive in.

Top US Privacy Laws as of 2021

Currently, there is no US federal law regulating the collection and use of private consumer data. We look forward to the day when we can update this paragraph with a less ridiculous opening line, but we're not holding our breath. It appears unlikely that the legislative bodies of the US government will work together to pass one in the near future. There have been an extraordinary number of proposals and all have gone nowhere, as of the time of this writing.

Instead, in the grand tradition of Colonial quilt-making, the US is steadily stitching together a piebald patchwork of State-level privacy laws, which we'll detail here. These laws protect the privacy of the residents within their states. Regardless of where your business is located, if you have customers living within a state with privacy legislation, you must comply accordingly to protect those customers' data.

One last note before we jump into the comically similar acronyms below, even if a federal law arrives, it may not preempt these state laws. It may just set the floor, leaving the door wide open for individual states to add more onerous privacy provisions as they see fit. Case in point, this is how HIPAA works. So today's patchwork quilt of privacy regulation in the US is likely here to stay. Snuggle down and get cozy with it.

California Consumer Privacy Act (CCPA)

California was the first US state to build on the privacy guidelines of the EU's GDPR (see below) and apply them to protecting the rights of California residents. In 2018, voters overwhelmingly supported the enactment of the California Consumer Privacy Act (CCPA). The law went into effect January 1, 2020, and was updated and expanded by Proposition 24—aka, the California Privacy Rights Act (CPRA)—which will go into effect January 1, 2023. The focus of both is to grant individuals rights over the use of their data: the right to access it, the right to know how it will be used, and the right to delete it.

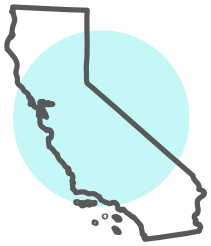

Companies must also give people an easy way to opt out of the sale of their personal information. Notably, the CPRA includes data on "households," not just individuals. This means companies must apply the same restrictions to information about a household and to data collectors that identify the income of a particular address.

Summary of the CCPA

> **Regardless of where your business is located, if you have customers living within a state with privacy legislation, you must comply accordingly to protect those customers' data.**

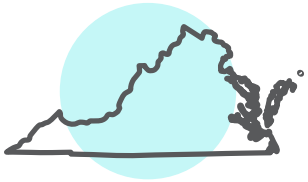

Virginia's Consumer Data Protection Act (CDPA)

This represents the prevailing alternate view on privacy in the US, considered more business-friendly and limited in scope. It was passed in 2021 and will go into effect in 2023. Like the CPRA, it allows consumers to opt out of the collection and sale of personal information. It also requires companies to provide adequate privacy notices, though "adequate" is undefined.

Essentially, this law states that companies must be transparent about what kind of information they are collecting and what they intend to do with it. If an individual doesn't want their data collected, they should be able to delete whatever the company has.

Summary of the CDPA

Colorado's Privacy Act (CPA)

Colorado's law is the nation's third, passed in 2021 and, like Virginia's, due to go into effect in 2023. The law looks like a mix of policies from CCPA and GDPR. It does have the opt-out of sale of personal information in CCPA but also uses phrases from GDPR, like data processors and data controllers.

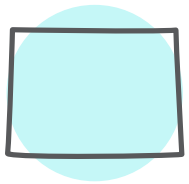

The Colorado law tends to be a bit more strict than California's in some ways, such as the requirement that companies prevent any "reasonably foreseeable risk" that a user's privacy could be compromised.

Summary of the CPA

Illinois Biometric Information Privacy Act (BIPA)

This is not a general data privacy law, but covers a specific sub-topic that is rapidly taking on greater significance. Although it passed in 2008, this law has been in the news recently for some massive class action lawsuits. In February 2021, plaintiffs in a BIPA class action suit against Facebook won $650 million in federal court as part of a proposed settlement.

BIPA states that companies must inform people in writing if their biometric information or identifiers will be collected and stored on a server. You must obtain the individual's written consent and have a policy that describes how you collect and use that information.

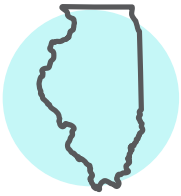

Imagine employers are taking fingerprint scans of employees so they can punch the time clock or using scans as an access point into the building. Many companies didn't publish a policy about what they do with that biometric info and didn't obtain consent. Attorneys have won multi-million dollar settlements against companies failing to comply with this law.

Under GDPR and CCPA, everything categorized as biometric info in the Illinois statute is considered PII. It is all treated as sensitive, including data on a person's genetic history, fingerprint to unlock a device, retina scan data, and gait analysis that identifies people based on how they walk. Given the bar-setting nature of these two laws, it's safe to expect this to be the case for the majority of data privacy legislation going forward.

Summary of BIPA

Health Insurance Portability and Accountability Act (HIPAA)

Whether you work at a healthtech company or you've been to a doctor's office in the past couple decades, you've probably crossed paths with HIPAA. The law first passed in 1996 with the objective of allowing workers to carry their healthcare insurance and healthcare rights between jobs.

Over the course of its passage through Congress and the decades since, its objectives have expanded to improve overall efficiency in healthcare, combat fraud, and ensure that an individual's protected health information (PHI) is adequately protected and kept private and confidential. Notably, in 2013, its reach extended from regulating how health plans and healthcare providers collect, process, and protect PHI to regulating how their business associates providing third party services handle any PHI they encounter, as well.

Comprehensively, and along with the advances promoted by the Health Information Technology for Economic and Clinical Health Act (HITECH), HIPAA has significantly streamlined how patient health information is collected and exchanged. It has also significantly impeded how healthtech companies are able to work with patient data to build their products.

The Rest of the US

"...over half of the US states have taken action to pass privacy legislation or establish a task force.

As of September 2021, over half of the US states have taken action to pass privacy legislation or establish a task force. Currently, there are bills in committee in Massachusetts, Minnesota, New York, North Carolina, Ohio, and Pennsylvania, several of which would be influential and are likely to pass in the near future. And in states where similar bills have been defeated, like Arizona, Connecticut, Texas, Florida, and Washington, it isn't necessarily because they aren't desired. Often they are defeated because they aren't deemed protective enough.

The International Association of Privacy Professionals (IAPP) site houses the latest updates on which states are working on what legislation regarding data privacy:

IAPP US State Privacy Legislation Tracker

Top International Privacy Laws as of 2021

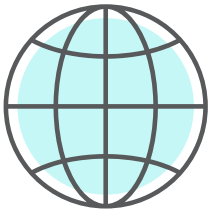

Globally, there are a large number of data privacy laws that impact international business and the movement of data across borders. The UN lists 128 countries that have passed legislation on the collection, use, and sharing of personal information to third parties without notice or consent. Reading them all is a project for your Data Privacy Officer.

UN Data Protection and Privacy Legislation Worldwide

For the purposes of your dev team, here are the top 4 (3 ½ really) international data privacy laws you need to know.

The EU's General Data Protection Regulation (GDPR)

This is the one that started it all. GDPR was passed by the EU in 2016 and went into force in 2018. It applies to any organization (private and public, for-profit and not) collecting data in Europe, with the exception of the UK, which is covered by its own Data Protection Act (see below).

Even before GDPR went into effect, there was a good deal of panic about what this would do to innovation. This intensified after Google was hit with a $57 million GDPR fine in 2019. That's when companies got serious about compliance, locking down personal information like phone numbers, email addresses, social security info, driver's license numbers, and financial data.

GDPR laws apply to virtually all global companies, even to those just doing development on a worldwide platform. This is known as its "extra-territorial effect." Even small US companies need to be aware of these data privacy guidelines if they are developing something that could be used by individuals in the EU.

Contrary to popular belief, GDPR does not require businesses to **obtain consent** from people before using their personal information for business purposes. It's just a popular approach because it opens the door pretty wide to allowing the data to be used as needed. But it's only one of six ways a business can identify and establish a legal basis for data processing. Briefly, here are the six permissible legal bases:

1. **You obtain consent from the data subject**
2. **You need to process the data to satisfy contractual obligations for the data subject**
3. **You need to process the data to comply with a legal obligation**
4. **You need to process the data to save somebody's life**
5. **You need to process the data to perform a task in the public interest or carry out an official function**
6. **You have a legitimate interest to process the data**

Yes, that last one is attractively vague. But before you go hog wild with your "legitimate" interests, keep in mind that the "fundamental rights and freedoms of the data subject" always override a business's interests. You would have to provide a pretty compelling case for your legitimate interest, and once established, it cannot change over time.

Summary of the EU's GDPR

The UK's Data Protection Act (DPA)

After the UK left the EU, they published their own version of GDPR which impacts all organizations doing business in the UK. It is very similar to GDPR but with a few additions. Most notably, there are more reasons given where it is permissible for companies to process the personal data of UK subjects, such as:

- employment, social security, and social protection
- health and social care
- archiving, research, and statistical purposes
- criminal convictions

But for the most part, you can consider it a post-Brexit copy/paste to salvage advances in data privacy.

Summary of the UK's DPA

Brazil's Lei Geral de Proteção de Dados Pessoais (LGPD)

Arguably the most beautifully named of the bunch, Brazil's law translates to the General Personal Data Protection Law (GPDL) and went into force in mid-2021. Brazil houses the largest population in Latin America and the 2nd biggest economy in the Western Hemisphere after the US, so LGPD stands to have a huge impact on business in the Americas.

While generally more lenient than GDPR, Brazil's version offers some additional consumer protections, such as the ability to request information about anyone who has received your private data.

Summary of Brazil's LGPD

Canada's Consumer Privacy Protection Act (CPPA)

This only counts as half a law because it is still being drafted at the time of this ebook's publication. It will expand on Canada's Personal Information Protection and Electronic Documents Act (PIPEDA), which covers what private sector firms that interact with Canadian residents must do to collect, use, and disclose private personal information. Notably, PIPEDA laws do not apply to nonprofits, political parties, associations, educational institutions, or hospitals.

CPPA will update this law to tighten accountability and control guidelines, as well as introduce more serious penalties for those who fall outside compliance.

Summary of the CPPA

Common Themes Across Privacy Laws

All of the above data privacy and consumer protection laws tend to share a few basic elements.

1. The **definition of personally identifiable information** (PII) is incredibly broad and getting broader. It's not just sensitive information, like address, credit card, social security, etc. It could be as simple as just name and email address, which is enough to obtain account access by identity thieves. This new definition represents a sea change for many in the US who have had a much narrower view of PII. Today, PII means any information that could identify, be used to contact, or be associated with a natural person, directly or indirectly.

> **Today, PII means any information that could identify, be used to contact, or be associated with a natural person, directly or indirectly.**

2. People have a **right to access, correct, and delete their data**. If you collect data, you are legally responsible for the handling of any data under your control. That data should be as accurate and complete as possible. Anyone who provides data must have access to their data and have a way to easily review the accuracy or completeness of that data for correction or deletion.

3. People have a **right to specify how their data is used**. If you are sharing that data with a third party, you must disclose who will receive the data and any possible risks to the individuals who provided the data. In some cases, individuals may opt out of the sale of their data. Companies collecting data should publish a policy document that clearly states their data collection purposes and how they will use the data.

4. Your organization should **assign a data privacy officer** with the role of managing data privacy compliance. The public should be able to reach this officer to report possible compliance violations. The officer ensures: that you only collect the personal data you absolutely need for the purposes stated in your data privacy document; that personal data stored on your servers is complete and accurate; and that data requires written consent before being collected, used, or disclosed.

5. You must protect the privacy of personal data with **security protocols that match the sensitivity** of the data. Breaches must be reported immediately to the compliance agency and to any individuals affected by the breach. Records of the breach must be made publicly available.

6. You must prioritize methodologies to **anonymize or de-identify PII**. Using de-identified or synthetic data significantly reduces your exposure under these laws in most circumstances. When it comes to software development, de-identified data enables you to work without concerns over data privacy compliance violations.

Core Concepts of Data Protection

Data Minimization

Most data privacy laws have a preference for using the least amount of personal information necessary to accomplish your purpose. The struggle now is between companies with the desire to monetize data versus the public's need to control how their personal information is used. Most privacy laws, like GDPR, operate under the assumption that you will not collect or hold personal information that is not required for the specified purpose.

If you are sending out an email newsletter, all you need to collect is the user's name and email address. If you collect address, phone number, and other data points, you're not practicing the idea of data minimization.

Developers typically need to have access to as much as possible to run effective testing, but that's the opposite of what these privacy laws advise. Collect the least amount of information necessary to achieve the operational purpose.

Data Privacy as a Human Right (DPaaHR)

GDPR specifically states that an individual's right to privacy with their personal information is a fundamental human right. That also means data control is a human right, which in part speaks to the skepticism felt toward machine learning, algorithms, and AI. The fear is that AI does not prioritize privacy and may engage in automated decision-making that can be biased or unnecessarily impede the rights of privacy.

DPaaHR is also part of the new Colorado law. In the US, there is a tendency to think, "I'm going to collect all this personal information and then it's mine. I own the data." These laws say, "No, the person identified in the data owns it. They control how the data is used and at any time they can modify it, delete it, or tell you not to process it in a certain way."

Data Localization

The EU and countries like Brazil have a preference for keeping data within their borders. Essentially, they are saying "We want you to keep that personal information in the country from which it originated because we know that we have laws that protect that information."

When data doesn't stay local, there are privacy compliance obligations that go along with the transfer process, including contracts, security requirements, privacy requirements, etc. That's why many companies encounter roadblocks in moving European data out of the EU and ultimately opt to avoid the complications of cross-border data restrictions.

The EU created a Privacy Shield Framework to handle data transfers with the US, but in 2020, the EU's Court of Justice handed down the Schrems II decision, invalidating the Privacy Shield Framework for being inadequate for the purposes of privacy protection.

Who Should Comply with These Laws?

The short answer here is: everyone. Ok, fine: everyone who collects and/or processes personally identifiable information. Which, by our latest count, is a whole lot of you. Regardless of where your company or team is based, you are obliged to comply with these laws if you do business with the residents of the regions in which the laws are in force.

And if there isn't a law in force in the region in which you do business, expect to see one soon. The momentum for expanded legislation is only growing stronger, both in and outside of the US.

But for those who want the cold, hard facts, here are a few quick litmus tests to help you determine if you need to comply **right now.**

> **Regardless of where your company or team is based, you are obliged to comply with these laws if you do business with the residents of the regions in which the laws are in force.**

If you check off any of these boxes, you need to comply with the related law **right now**

For CCPA (California)

- ☐ You collect or process data on a minimum of 50,000 California residents
- ☐ Your business is based in California and brings in over $25 million in revenue

For HIPAA (USA)

- ☐ You are a health plan provider
- ☐ You are a healthcare clearinghouse
- ☐ You are a healthcare provider who conducts financial and administrative transactions electronically
- ☐ You are a business associate of any of the above, providing third party services and activities during the course of which you will encounter PHI (hi there, healthtechs! 👋)

For GDPR (EU)

- ☐ You offer products or services to individuals located in the EU
- ☐ You have physical locations or salespeople in the EU

For DPA (UK)

- ☐ You offer products or services to individuals located in the UK
- ☐ You have physical locations or salespeople in the UK

For LGPD (Brazil)

- ☐ You collect or process data from individuals located in Brazil
- ☐ You offer products or services to individuals located in Brazil
- ☐ You have physical locations or salespeople in Brazil

For those of you looking to get your ducks in order, if you check off any of these boxes, you will need to comply with the related law when it goes into effect in 2023.

For CDPA (Virginia)

- ☐ You conduct business in Virginia, OR
- ☐ You produce and sell products that target Virginia residents

AND

- ☐ You control or process the personal data of at least 100,000 Virginia consumers during a calendar year, OR
- ☐ You control or process the personal data of at least 25,000 consumers and derive over 50% of gross revenue from the sale of personal data

For CPA (Colorado)

- ☐ You conduct business in Colorado, OR
- ☐ You produce or deliver commercial products or services that target Colorado residents

AND

- ☐ You control or processes the personal data of 100,000 consumers or more during a calendar year, OR
- ☐ You control or process the personal data of at least 25,000 consumers and derive revenue or receive a discount on the price of goods or services from the sale of personal data

If you check off any of these boxes, you need to comply with the related law by 2023

Is anyone exempt?

Yes.

- Non-profit organizations and government agencies are exempt from CCPA and CDPA. But they are not exempt from CPA, GDPR, DPA, or LGPD.
- Entities governed by GBLA are exempt from CDPA and CPA.
- Entities regulated by HIPAA are exempt from CDPA.
- The PHI controlled or processed by entities regulated by HIPAA are exempt from CCPA, but the PII controlled or processed by entities regulated by HIPAA are not exempt from CCPA.
- Everyone is exempt from CPPPO because it doesn't exist... yet. (Help us, Obi-wan Connecticut. You're our only hope.)

Adjust your check boxes accordingly.

Penalties

GDPR touched off a firestorm with its massive penalties. A story in the Wall Street Journal reported that Amazon has racked up $425 million in fines under GDPR.

Amazon faces potential $425 million EU privacy fine

That may be just a drop in the bucket for Amazon, but it could devastate a smaller company.

Fines can be up to 2 - 4% of worldwide revenue, so multi-billion dollar companies are in the greatest jeopardy of the largest fines. What the EU is really going after is companies that expose private data irresponsibly through data breaches and don't comply with the principles behind the law.

The penalties for violations to California's laws are more rare and lenient for now. Under CCPA, if you don't follow the regulations, it can be $2,500 per negligent violation or $7,500 per intentional violation. But if you receive a notice from the Attorney General of California who enforces CCPA finds that you're not doing something right, you do have a 30 day period to cure.

> **Starting in 2023, the 30 day period to cure is going away, and instead, you will be fined immediately.**

Currently, the chances of getting fined are relatively low under CCPA but that will change after 2023. The 30 day period to cure is going away, and instead, you will be fined immediately. Combine this with the fact that the Virginia and Colorado laws go into effect in 2023, as well, and it's clear to see that the time to get your house in order is now.

The last point to watch out for is data breaches. Under CCPA, there's a specific provision that says if you fail to use reasonable security measures to secure personal information, a plaintiff can recover up to $750 per consumer per incident. So if you had 30,000 people whose information was compromised at $750 per person, that could be a lot of money.

Data Privacy Compliance Action Plan

Or as we like to call it, a DPCAP for CCPA, DPA, CDPA, CPA, CPPA, and whatever other combination of Cs, Ds, Ps, and As the world of legislation can throw at us. Whether you're just getting started on your journey to compliance, or you're looking to refresh and strengthen existing processes, here's a solid approach to data governance that will set you up for success across the future landscape of data privacy regulations.

1

Step 1: Data Mapping

Your first step is to understand the flow patterns of where data is coming from, where it is going, and what data controls you have in place. In other words, show your Ops teams some love, and give their data pipeline documentation the attention it deserves.

2

Step 2: Gap Analysis

Once you've mapped your data, it's time to document how your team is currently handling that data and how each step of the way measures up against the privacy laws that apply to your company.

3

Step 3: Build Your Privacy Framework

Use everything you've gathered and learned to build out a privacy framework. Consider this your first line of defense for maintaining compliance and avoiding violations down the line. Your privacy framework should address all of the common themes in data privacy listed above, to define how data must be collected and processed by your organization.

Your team can then implement the framework company-wide to close the gaps you identified earlier and forge ahead with their work, confident that they aren't violating any individual's privacy or any country's privacy law. (Sounds nice, right?)

When it comes to software development, your privacy framework will likely get the biggest bang for its buck by adhering to the practice of data minimization. Developers don't need real data to build and test their products; they just need realistic data. Restricting your dev team's access to production data significantly lowers the risks of data breaches and accidental exposure while alleviating cross border data issues, as well.

Alright, great, you've removed all prod data from staging. Those test databases are empty—clean as a whistle. Have fun, developers!

Jk. We've got a fix for you that'll solve your compliance woes in a single shot.

How to Make Your Dev Team Compliant

When it comes to data privacy compliance in software testing and development, there is a silver bullet: **data de-identification.** Each of the laws currently in force stipulate that adequately de-identified data are no longer subject to their regulations and restrictions. In other words, if you de-identify your data to their standards, you can do what you want with it, like send it to that off-shore development team that's never had a decent test database to shake a line of code at. So what are those standards?

The requirements for data de-identification under CCPA are a good yardstick for all the other privacy laws. CCPA says that data has been de-identified when "it cannot reasonably identify, relate to, describe, be capable of being associated with, or be linked directly or indirectly" to a specific individual. #vaguelawing. "Cannot reasonably identify" is very open-ended, so let's make that statement more concrete with a technical term:
your data is de-identified when it cannot be re-identified.

This may sound like a straightforward no-brainer, but just think of all the times "anonymized" data has been traced back to real individuals. To protect yourself against these data fiascos, your de-identification process should include technical safeguards and business policies that effectively eliminate the risk of re-identification. Your to-dos include:

> **"Each of the laws currently in force stipulate that adequately de-identified data are no longer subject to their regulations and restrictions.**

1

Performing de-identification in a way that provides mathematical guarantees of data privacy

2

Ensuring that no one on your team will attempt to re-identify the data

3

Preventing inadvertent release of the data, even if it is de-identified data

In doing all this, you'll want to make sure your approach is valid and robust enough to scale and serve the needs of the enterprise. It can be beneficial to bring in an outside expert in data privacy to avoid bigger issues down the line. Here are some best practices to keep in mind.

Differential Privacy

Differential privacy is a property that can be applied to mathematical algorithms for data synthesis that adds statistical noise to datasets to significantly reduce the ability to re-identify an individual person in a given dataset. It provides a mathematical guarantee against data re-identification, hardening the security of synthetic outputs.

For example, an algorithm for generating data might perform some statistical analysis on a given dataset (such as the dataset's mean, variance, mode, etc.). The algorithm is then said to be "differentially private" if, upon looking at the synthetic output it generates, it is impossible to tell whether a given record from the original dataset was included in the calculation, regardless of how typical (or atypical) that record is. Essentially, differential privacy is a way of guaranteeing an individual's privacy when working with a larger dataset in aggregate.

If you make your data generation algorithms differentially private, you are providing mathematical guarantees against data re-identification.

Privacy by Design

Thinking about the software development process at a higher level, during the design phase, privacy by design means building in the tools to ensure privacy in your processes from Day 1. If you collect personal information through your software, plan for the data controls you'll need to have in place. Make sure your users have the ability to access their data, export their data, and delete their data in a simple, direct way.

Audit logs are also essential so that you can immediately see who has accessed which pieces of personal data and when. Encrypting PII is a great idea, both at the storage level and at the field level. The goal is to build in plenty of privacy controls and features in the underlying product so that you can meet the requirements of any privacy laws that may be on the horizon.

Another valuable strategy is to tag fields when they contain personal or sensitive information. This will help companies identify where the greatest privacy risks are in their software.

Lastly, you should assign role-based access and need-to-know controls on the data being used in development. Not everybody needs access to production data in order to safely de-identify it before it makes its way into your lower environments. The tools you use to de-identify your data should support this.

A Special Shout-out to the Healthtechs

De-identification under HIPAA is a little more complicated than your run-of-the-mill data privacy laws because even after removing certain identifying information you may still end up with a designated record set or limited dataset that's subject to HIPAA. If so, you will still have to sign HIPAA business associate agreements and comply with all of the requirements under HIPAA.

To achieve compliance under HIPAA, you need to ensure that you're working with data de-identified to the law's specific standards. There are two ways to do this:

1

Have a subject matter expert with statistical and scientific experience apply principles and methods to your de-identification process, to determine that the risk of re-identification is small. They will have to document the methods and results of the analysis that justify their determination.

2

Remove all identifiers, including name, geographic info, dates, address, telephone number, email, SSN, medical records, URLs, IP addresses, biometric identifiers, images, and "any other unique identifying number, characteristic, or code."

Entities in the HIPAA space tend to prefer option 1 because they want the comfort of having a report and a document that specifies the methodology and confirms that the data can't be re-identified. Not all companies can get an expert to do that so that's why option 2 exists.

Wait....

...Why Should I Care Again? Aka, The Takeaways

It's simple.

1. Data privacy laws are here, getting stronger, and not going away.
2. Data de-identification is all your dev team needs to comply.
3. The sooner you build out your privacy framework, the fewer headaches you'll have down the road.

All of these laws effectively set the same data privacy standards:

- The **definition of personal identifiable information** (PII) is incredibly broad and getting broader.
- People own their data. They have the **right to access, correct, and delete** it from your database.
- People have a **right to specify how their data is used.**
- Personal data should be protected with **security protocols that match the sensitivity** of the data.
- Prioritization should be given to **de-identifying PII.**

Working with de-identified and synthetic data in dev and testing environments simplifies developers' responsibilities on several fronts:

1. It minimizes data exposure, lowering the risk of a breach or leak.
2. It eliminates the complications of working with data on virtual desktops.
3. It reduces the amount of due diligence required when working with vendors.
4. It circumvents the roadblocks encountered in cross border data flows.
5. It allows security to focus on ensuring privacy compliance in one environment only: prod.

Now that you're convinced, you may be wondering *how* to de-identify your data. Luckily, we've got an ebook on that too. Head on over to Test Data 101.

TONIC

THE FAKE DATA COMPANY

Tonic.ai empowers developers while protecting customer privacy by enabling companies to create safe, synthetic versions of their data for use in software development and testing. Founded in 2018, with offices in San Francisco, Atlanta, and New York, the company is pioneering enterprise tools for database subsetting, de-identification, and synthesis. Thousands of developers use data generated with Tonic.ai on a daily basis to build their products faster in industries as wide ranging as healthcare, financial services, logistics, edtech, and e-commerce. Working with customers like eBay, The Motley Fool, Flexport, and Everlywell, Tonic.ai innovates to advance their goal of advocating for the privacy of individuals while enabling companies to do their best work.

Chiara Colombi

Product Marketing Manager

TONIC

A bilingual wordsmith dedicated to the art of engineering with words, Chiara has over a decade of experience supporting corporate communications at international companies, including Finmeccanica, Enel, and RAI. She once translated for the Pope; it has more overlap with translating for developers than you might think.

Omed Habib

VP of Marketing

TONIC

A developer at heart, Omed fell in love with fake data as a way to improve developer productivity. He formerly led Product Marketing teams at AppDynamics, Harness.io, and helped launch startups from inception to unicorns. When not faking data, Omed keeps busy geeking out on all things tech, photography, and cooking.