KASOWITZ

Data Privacy, AI Regulatory, and Compliance Update: 2026

2026 will be one of the more active periods of privacy compliance in recent memory. Multiple U.S. state laws covering comprehensive privacy, data brokerage, and age verification will take effect, alongside major international developments: Vietnam's first national data protection law goes into effect, the United Kingdom is rolling out digital verification services, Australia has new transparency mandates for automated decisions, and the European Union's Data Act design obligations all take effect. Collectively, these measures mark a global shift from privacy "as disclosure" to privacy "as infrastructure," where technical design, interoperability, and verified user control define compliance as much as policy and consent.

2026 also marks a shift in global AI governance. Across the United States, Europe, and Asia, jurisdictions are implementing the first binding regulatory regimes designed to move AI oversight from principle to enforceable obligation. While the EU AI Act's high-risk system rules likely take effect in August 2026 (pending potential delays proposed by the EU Commission), several U.S. states, particularly California, Texas, and Colorado, are also entering the compliance phase of their AI and data-privacy programs. Together, these developments may signal the end of the AI "self-regulation" era and the rise of multi-layered, legally mandated governance frameworks.

Please read below for information on new laws and amendments taking effect in early 2026.

<u>Three New U.S. Comprehensive Privacy Laws and CCPA Regulations</u>

Indiana, Kentucky, and Rhode Island transition from planning to enforcement on New Year's Day, 2026.

Indiana's Consumer Data Protection Act (ICDPA) largely follows the "Virginia" model (consumer rights, controller/processor duties, DPIAs for certain high-risk processing), and provides standard data-level and entity-level exemptions. The ICDPA applies if you control/process personal data of 100,000+ Indiana consumers in a year, or 25,000+ and derive >50% of gross revenue from selling personal data. Notably, DPIA obligations apply to activities created or generated after December 31, 2025, and are not retroactive to any activities before January 1, 2026. In addition, Indiana follows the minority of states in defining the "sale" of personal data as an exchange for monetary consideration rather than the broader "valuable consideration" used in states like California. Last, there is a mandatory 30-day cure period before enforcement from the AG may commence.

Kentucky's Consumer Data Protection Act (KCDPA) follows many of the same baseline requirements as Indiana and other Virginia model state privacy laws, but does not provide an exemption for aggregated data, nor does it have a universal opt-out mechanism requirement. The KCDPA applies if you process 100,000+ Kentucky consumers, or 25,000+ and >50% revenue from selling personal data. Kentucky also diverges from other state laws by providing entity-level exemptions for certain niches, including organizations involved in assisting law enforcement agencies with insurance-related fraud, small telephone utilities, Tier III Commercial Mobile Radio Service providers, and municipal utilities that do not sell or share personal data with third-party processors.

Rhode Island's Data Transparency and Privacy Protection Act applies to entities doing business in RI or targeting RI residents, if they processed 35,000+ residents' data in the prior year, or 10,000+ and >20% revenue from data sales—lower thresholds than many states. Core obligations include standard consumer rights, controller duties, processor contracts, and enforcement. But there are two notable policy choices: no mandated UOOM honoring (i.e., no prescriptive requirement to accept browser/global signals as in CA/CO), and no statutory cure period, raising immediate enforcement risk on day one.

Last, the CCPA's cyber, risk, and automated decision-making technology (ADMT) regulations become effective on January 1, 2026, with delayed deadlines for most requirements.

Texas SB 2420 ("App Store Accountability Act") and Texas CUBI Amendments (via HB 149)

Texas's App Store Accountability Act will require app stores to verify age at account creation with commercially reasonable methods, place minors under a supervising account, and obtain and pass parental-consent status downstream to developers before downloads or in-app purchases. Practically, app stores will likely need to integrate platform signals with lawful bases (verifiable parental consent), minimization (what attributes are actually stored), retention (delete on account status change), security (access controls, audit logs), and downstream vendor controls (SDKs can't exceed the scope of consent). Platforms are shipping developer-facing changes that underscore privacy tradeoffs, while publicly flagging privacy risks from ID collection for routine app downloads.

Texas's HB 149 (TRAIGA) also introduces amendments to its Capture or Use of Biometric Identifier Act codified at Bus. & Com. Code Section 503.001, clarifying that a person is not informed and has not consented to biometric capture/storage solely because an image or other media containing their biometric identifiers is available on the internet or other public sources unless the individual themselves made the media public. This will effectively close the "publicly available image equals implied consent" argument for commercial biometric use in Texas. CUBI also includes pre-collection consent and notice requirements, security requirements, and record retention obligations related to biometric identifiers.

Vietnam's PDPL (Law No. 91/2025/QH15)

Vietnam's first statutory personal data protection law replaces the more limited Decree 13 and realigns Vietnam's data privacy regime.

What makes the PDPL especially consequential is its extraterritorial scope. Any foreign entity processing personal data of Vietnamese citizens or individuals of Vietnamese origin residing in Vietnam potentially falls within the scope, even absent a physical presence in the country. For cloud services, analytics platforms, or marketing database providers, that means that Vietnamese data cannot be ignored simply because there's no local office. Rather, you must treat it as a regulatory domain upfront.

Under the PDPL, consent also becomes more granular and more demanding. The default assumption is that no personal data may be processed without explicit, specific consent tied to a declared purpose. Blanket consents or vague, catch-all permissions will no longer pass muster. Moreover, the PDPL requires documentation of consent, retention of consent records, clarity on withdrawal pathways, and, in certain high-impact scenarios, prior approval or registration of the consent framework itself. For companies reliant on data aggregation, profiling, or segmentation, this may be a fundamental governance and operationalization shift.

Furthermore, cross-border data transfer obligations under the PDPL will become more complex. Vietnam is not an EU-adequate jurisdiction, so cross-border transfer controls and impact assessments/approvals may be a serious operational hurdle for global data enterprises, and companies may need to lean on contractual tools, binding rules, or even localization fallback strategies.

California DELETE Act Registration

California's data broker registration requirements under the DELETE Act require qualifying entities to register with the CPPA by January 31, 2026. The statutory registration fee is $6,600 per year, and failure to register triggers a fine of $200 per day. Among other things, registration requires providing statistics and information about applicable exemptions to the CPPA, which will require in-scope organizations to do more than simply sign a check before meeting compliance requirements.

U.S. State-Level AI Laws, Trump's EO, & Vietnam

California Generative AI Transparency Act (AB 2013)

Signed September 28, 2024, and effective January 1, 2026, California AB 2013 requires qualifying developers of publicly available generative AI systems to disclose information about the data used to train their models. Covered providers, including entities that substantially modify the generative AI Systems and make it accessible to California users, must post a publicly available summary describing the categories and sources of training data. This requirement applies regardless of company size or profit status. This regulation highlights the importance for businesses to understand and be transparent about the data used to train their generative AI systems at the outset. While the law contains no private right of action, the California Attorney General may treat non-compliance as an unfair or deceptive business practice under the UCL, carrying civil penalties of up to $2,500 per violation.

California Frontier AI Safety Act (SB 53)

Also effective January 1, 2026, California SB 53, the Transparency in Frontier Artificial Intelligence Act, targets so-called “frontier AI models” trained with compute exceeding 10^{26} FLOPs. It will require large developers (annual revenue >$500 million) to publish Frontier AI Safety Frameworks detailing how catastrophic risks are identified, mitigated, and monitored, and to issue transparency reports before deploying any new frontier model. Developers must also report “critical safety incidents” to the California Office of Emergency Services within 15 days (24 hours for imminent threats). Additionally, applicable businesses need to create anonymous whistleblower channels that provide intake systems for employees who raise concerns about catastrophic risks. Civil penalties reach $1 million per violation.

California CCPA ADMT Regulations

California’s Automated Decision-Making Technology (ADMT) Regulations, finalized by the California Privacy Protection Agency (CPPA) in July 2025, mark the state’s first attempt to regulate algorithmic decision systems under the CCPA/CPRA framework. Although formally effective January 1, 2026, the ADMT-specific compliance date is deferred to January 1, 2027, giving businesses a one-year runway to implement notice, opt-out, and risk-assessment processes.

The regulations define ADMT as any technology that processes personal information and replaces or substantially replaces human decision-making. This language narrows earlier drafts that would have covered routine automation, confining the rules to systems that issue or materially determine outcomes without human intervention. They apply when ADMT is used to make “significant decisions” producing legal or similarly significant effects, such as those affecting employment, credit, housing, healthcare, education, or essential services. Assistive tools that merely aid human judgment generally fall outside scope.

Businesses using ADMT for significant decisions must provide pre-use notice, offering consumers plain-language disclosure of the tool’s purpose, logic, and potential impacts. Consumers gain an explicit opt-out right, and those subject to ADMT decisions may request access and explanations describing the main factors that influenced the outcome. Before deployment, companies must perform and document privacy-risk assessments evaluating foreseeable harms, safeguards, and mitigation measures. Larger entities, particularly those exceeding $100 million in revenue, must also conduct independent cybersecurity audits and provide annual attestations to the CPPA beginning in April 2028.

Texas Responsible AI Governance Act (TRAIGA, HB 149)

When Texas enacted HB 149, the Texas Responsible Artificial Intelligence Governance Act (TRAIGA), in June 2025, it aimed to balance innovation with restraint. Taking effect January 1, 2026, TRAIGA is often mischaracterized as a private-sector AI law; in fact, its heaviest mandates fall on government agencies, while private businesses face only targeted prohibitions. Still, those prohibitions are legally enforceable and carry significant reputational and civil-penalty risk.

The statute's jurisdictional reach is broad: any person developing, deploying, or distributing an AI system in Texas, or offering AI-related products or services consumed by Texas residents, comes within scope. However, private entities face no new transparency duty. Their obligations arise instead from TRAIGA's catalogue of prohibited AI uses: intentionally deploying or designing AI to manipulate users into self-harm, violence, or criminal acts; using AI to violate constitutional or statutory rights; generating unlawful or pornographic content (particularly involving minors); or employing deceptive trade practices to induce harmful behavior.

TRAIGA also amends Texas's biometric-identifier statute (Bus. & Com. Code § 503) to clarify that private developers may use biometric data for training, processing, or development, provided the AI system is *not* intended to uniquely identify specific individuals. This change eases constraints for model-training activities while preserving consent rules for identification uses. In healthcare settings, however, both public and private providers must disclose AI use at or before the point of care (except during emergencies). These targeted duties make Texas's framework more permissive than Colorado's, which mandates risk assessments and algorithmic-discrimination safeguards across industries.

In terms of enforcement, fines for violations range from $10,000 for minor or curable offenses up to $200,000 per uncurable violation, with daily penalties for continuing misconduct; professional licensees (e.g., physicians, attorneys) risk additional sanctions up to $100,000 and potential suspension.

Vietnam's Law on Artificial Intelligence

On December 10, 2025, the National Assembly of Vietnam passed the Law on Artificial Intelligence, formally separating AI governance from the broader PDPL. The law goes into effect March 1, 2026, giving global companies less than 3 months to prepare for implementation. Three key takeaways include:

1) The "Deployer Pays" Liability Shift: Unlike frameworks that heavily burden AI model developers, Article 29 of the law places the financial bullseye on the implementer (deployer). If a high-risk system causes damage (even if you fully complied with all regulations), you are likely liable to compensate the victim. You can seek reimbursement from the vendor later, but only if your contract allows it.
2) Agile "High-Risk" Definitions: The specific list of "high-risk" systems is not hard-coded in the law. The Prime Minister has the authority to update this list without legislative amendment. This creates a regulatory environment that can pivot quickly, but also brings some uncertainty.
3) Extraterritorial Scope: The law applies to foreign organizations participating in AI activities in Vietnam. If you are deploying models that interact with Vietnamese users or data, you are likely in scope.

White House Executive Order "Ensuring a National Policy Framework for Artificial Intelligence"

On December 11, 2025, the Trump administration issued an Executive Order purporting to move towards a federal standard on AI. The EO gives the Attorney General 30 days to establish a task force to challenge state AI laws that conflict with this new national policy, impede interstate commerce, or are deemed unconstitutional, so by the end of January 2026, we should have additional guidance.

Moving Forward

By the end of 2026, privacy compliance may look less like a checklist and more like a true ecosystem of interconnected parts. This year's wave of new laws, from U.S. state regimes and Texas's biometric and broker amendments to the EU Data Act, Vietnam's PDPL, and Australia's ADM transparency rule, will push organizations to treat privacy as a design and infrastructure problem, not just a paperwork and documentation exercise.

The common thread is accountability: regulators are embedding privacy expectations into the architecture of consent, identity, and data portability itself. For companies operating across borders, 2026 is the year to align governance, engineering, and policy into a single privacy operating model. Those who adapt early, building systems that are interoperable, explainable, and accountable will find compliance not just easier, but strategic.

* * *

Kasowitz's Data Strategy, Privacy, and Security team has deep knowledge in the data, privacy, and security sectors, and is familiar with the potentially existential risks faced by companies that rely on data as an engine of commerce and innovation. Global data, AI, privacy, and security threats are "bet the company" issues that Kasowitz is well equipped to handle. Our team consists of seasoned lawyers who have worked at or represented the largest and most innovative companies in the world, former regulators, and former government attorneys. We leverage our extensive subject matter knowledge to support companies through global privacy and technology counseling, regulatory support in the AI, privacy and security space, litigation, and incident preparedness and response.

For more information, please contact:

Brandy Worden
Partner
bworden@kasowitz.com

Frederick C. Bingham
Associate
fbingham@kasowitz.com

Cyrus Borhani
Associate
cborhani@kasowitz.com