2026 AI Laws Update: Key Regulations and Practical Guidance

Insights
February 5, 2026

Overview

This client alert provides a high-level overview of key AI laws enacted or taking effect in 2026. With President Trump's December 2025 Executive Order signaling federal intent to consolidate AI oversight, new comprehensive governance frameworks in Colorado and California, and evolving international requirements under the EU AI Act, companies developing or deploying AI systems face a rapidly shifting compliance landscape.

For ease of use, this alert is organized by topic, and the links below allow you to navigate directly to sections.

National AI Regulation

- **U.S. Executive Order on "Ensuring a National Policy Framework for AI":** On December 11, 2025, President Trump signed an Executive Order titled "Ensuring a National Policy Framework for Artificial Intelligence" (the "EO"). The EO signals an intent to consolidate AI oversight at the federal level, counter the expanding patchwork of state AI rules, and maintain U.S. global AI dominance by: (1) discouraging states' authority to regulate AI through litigation, spending, and standard-setting levers; (2) targeting state regulatory action in areas such as algorithmic transparency, bias mitigation, and regulation of high-risk AI uses; and

(3) promoting minimally burdensome national AI standards. Implementation of the EO is likely to face significant legal and political challenges. Notably, the EO maintains that federal efforts will not preempt state authority in areas such as child safety, AI infrastructure, and governmental AI procurement. The EO itself does not preempt, suspend, or invalidate current and enacted state AI laws and further federal action is required. Companies should continue efforts to comply with existing state laws until courts and agencies clarify the EO's reach.

- **International AI Regulations:** Outside the U.S., many governments have released AI regulations, some of which will have global effect and applicability. For example, the European Union (the "EU") has adopted binding legal frameworks that extend to non-EU based organizations, most notably the EU AI Act, which imposes significant obligations on high-risk and general-purpose AI models. These obligations will come into force over the next few years on a staggered basis, and include controls on data quality, transparency, human oversight, and monitoring discrimination. The EU Data Act adds data-sharing obligations. Additionally, under the GDPR's Article 22, individuals, including job applicants and employees, already have the right to not be subject to fully automated decisions (e.g., hiring, promotions and performance reviews) that significantly affect them, unless specific legal safeguards are in place.

State AI Regulation

- **Comprehensive AI Governance:** Several states have enacted or finalized broad AI governance statutes that impose affirmative risk management, documentation, and oversight obligations for certain high-impact AI systems, with enforcement beginning in late 2025 and 2026. While most startup companies will not meet statutory applicability thresholds, these laws are already shaping vendor contracting practices and downstream compliance expectations, particularly through AI-specific addenda and third-party risk allocation.

- **Consumer AI Interactions:** States have also begun targeting consumer-facing AI interactions (e.g., chatbots, AI companions, and algorithmic pricing based on consumer personal data) by requiring clear disclosures, safety protocols around high-risk use (e.g., self-harm or minors), and limits on use of personal data. Even for companies operating outside the most heavily regulated AI features, this regulatory trend may affect company choices regarding product design and contracting norms, including expectations for transparent AI labeling, crisis-response playbooks, and tighter representations and warranties around use of AI in consumer interactions.

- **AI Content Transparency:** States have begun to address AI content transparency by requiring developers, platforms, and advertisers to disclose when content is AI-generated, summarize AI training data, and display warning labels tied to AI-mediated or "addictive" experiences, particularly for young users. Companies are expected to make conscientious design decisions around provenance tooling, on-content AI labels, and risk-oriented warnings.

- **Automated Decision-Making Tools (ADMTs):** States and cities are treating ADMTs, including resume screeners, interviewing tools, HR tools used managing and evaluating talent, and other tools that "substantially assist or replace" human discretion, as an early beachhead for AI regulation. Emerging laws and bills in jurisdictions such as New York City, Colorado, Illinois, and New York State layer bias-audit, notice, recordkeeping, and human-review requirements onto these tools, and often require companies to allocate responsibility for compliance and algorithmic performance in their contracts with technology vendors.

- **Anti-Discrimination and Civil Rights:** Civil rights regulators are making clear that automated systems do not sit outside traditional anti-discrimination frameworks. Federal and state agencies, such as the EEOC, FTC, and state civil rights departments, have emphasized that existing employment, credit, housing, disability, and consumer protection laws apply equally to AI-mediated decisions, and that organizations can face liability for disparate impact, failure to accommodate, or unfair practices even when they rely on third-party models.

U.S. Executive Order on "Ensuring a National Policy Framework for AI

On December 11, 2025, the President issued an EO titled "Ensuring a National Policy Framework for Artificial Intelligence". The EO broadly calls for a national policy framework on AI and tasks U.S. agencies "to sustain and enhance U.S. global AI dominance through a minimally burdensome national policy framework for AI" by preempting state regulation of AI through federal lawsuits and withholding federal funds. While the EO does not immediately overrule existing state laws, it introduces mechanisms that could narrow or challenge the expanding patchwork of state-level AI obligations. Key takeaways include:

- **Streamlining AI governance at the federal level.** The EO aims to reduce multi-state compliance burdens that fall heavily on startups with lean compliance

functions. Federal agencies are directed to evaluate whether uniform federal standards should replace or supersede differing state requirements.

- **Federal oversight of state AI laws and federal funding as leverage.** A new litigation task force may challenge state regulations that federal policymakers view as "onerous," burdensome, or otherwise obstructing innovation. The EO directs agencies to evaluate whether federal grants may be conditioned on states aligning with the federal AI framework. While this does not impose direct compliance obligations on startups, it could materially change the regulatory landscape in innovation-heavy states.
- **Companies must closely monitor further developments.** The EO is likely to face substantial implementation and legal challenges, further complicating an already fragmented regulatory landscape that also includes evolving international AI regimes, particularly in the EU and UK. Companies should closely monitor agency actions implementing the EO, anticipated state resistance, and additional federal agency efforts to advance a "minimally burdensome national policy framework for AI."

The EO directs several agencies to establish enforcement procedures, as summarized in the table below:

EO Mandate	Details	Potential Impact
Evaluation of State AI Laws	***Within 30 days of the EO***, the Department of Justice and Attorney General ("AG") must establish an "AI Litigation Task Force" to identify and challenge state AI laws that the U.S. administration deems unconstitutional, unlawful, or preempted by federal policies and regulations. The Task Force is instructed to consult with senior White House advisors, including the Special Advisor for AI and Crypto, the Assistant to the President for Science and Technology, and the Assistant to the President for Economic Policy, to determine which state laws should be subject to legal challenge. ***Within 90 days of the EO***, the Secretary of Commerce, Special Advisor for AI and Crypto, Assistant to the President for Economic Policy, the Assistant to the President for Science and Technology, and Assistant to the President and Counsel to the President must identify state laws that: (1) are "onerous" and conflict with "minimally burdensome" federal policy and should be referred to the AI Litigation Task Force for potential legal challenge, and (2) promote AI innovation consistent with the aim of the EO (together, the "Evaluation"). In particular, the Evaluation must identify state laws that require disclosure or reporting compliance obligations that infringe about First Amendment rights (e.g., laws that "require AI models to alter their truthful outputs" or that otherwise "compel AI developers or deployers to disclose or report	Action by the AI Litigation Task Force may modify, limit, or invalidate existing state AI regulations. For example, the EO specifically cites the Colorado AI Act as a state law that bans algorithmic discrimination in a manner that compels AI models to produce false outputs. The Colorado AI Act will likely be subject to review by the AI Litigation Task Force this year. ***Companies must closely monitor any challenges of state AI laws by the AI Litigation Task Force. However, for the time being, companies should continue compliance efforts to meet existing state AI regulations.***

	information" in a manner that would violate constitutional rights).	
Federal AI Legislation and Preemption of State AI Laws	Within 90 days after publication of the Evaluation, the Federal Communications Commission ("FCC") must begin a process to determine whether to adopt a federal reporting and disclosure standard for AI models, which is intended to preempt conflicting state AI laws. Within 90 days of the EO, the Federal Trade Commission ("FTC") and Special Advisor for AI and Crypto will issue guidance clarifying when the FTC Act's prohibition on "unfair and deceptive acts or practices" applies to AI models. The policy statement must explain circumstances under which the FTC Act preempts contrary state AI laws that require alterations to "truthful" AI outputs. Additionally, the EO tasks the Special Advisor for AI and Crypto and the Assistant to the President for Science and Technology with preparing a legislative recommendation to establish a uniform federal AI framework that preempts state laws in conflict with the national policy. Note that the EO directs the legislative recommendation to exclude state AI laws relating to the following topics from proposed federal preemption: • Child safety protections, • AI compute and data center infrastructure, • State procurement and governmental use of AI, and • Other topics, as to be determined.	The scope of future federal AI regulation could broadly impact developers of frontier AI models, as well as downstream deployers and distributors of such AI models. Federal regulations could broadly implement baseline requirements to address algorithmic discrimination, bias audits and reporting, content disclosures and marketing practices, and other such consumer protection concerns. ***While the EO expresses the intent to streamline compliance requirements for companies using or distributing AI services nationwide, the EO does not establish new federal AI governance regulations and defers enforcement to applicable federal agencies. Companies must closely monitor agency initiatives to implement the EO.***
Federal Funding Eligibility Restrictions	***Within 90 days***, the Department of Commerce must issue a policy notice specifying when states remain eligible for federal infrastructure deployment (e.g., fiber installation) funding under the Broadband Equity, Access, and Deployment ("BEAD") program (the "Policy Notice"). States with "onerous" AI laws will be ineligible for BEAD funding. Further, states that are granted BEAD funding are restricted from using the federal program grant to finance supporting functions like planning, administration, outreach, research/data, or other such non-construction uses. ***Additionally***, the EO directs federal departments and agencies to review their discretionary grant programs with the Special Advisor for AI and Crypto to determine whether such grants can be conditioned on states agreeing not to enact or enforce AI laws that conflict with the EO's policy objectives. States with enacted AI laws may enter into a binding agreement with the relevant agency not to enforce any such state AI laws during the period in which the state receives discretionary federal funding.	State agencies and quasi-public entities may tighten up sub-grant and -award conditions to demonstrate alignment with federal funding eligibility requirements promulgated under the Policy Notice, with potential downstream impacts on AI companies operating in sectors that rely more heavily on federal grants (e.g., education, energy, finance, healthcare, etc.). ***Companies reliant on governmental funding will need to review the forthcoming Policy Notice, as well as monitor changes to federal and state eligibility requirements.***

PRACTICAL GUIDANCE

For companies using or creating AI-enabled technologies, the EO is not "an amnesty or moratorium [of state AI laws], but rather a statement of principles and set of tools" for the current administration to address "onerous and excessive" state AI laws. The EO does not establish any federal AI standards or regulations on its own and, absent further Congressional and federal agency actions, the EO merely signals federal intent to address and govern AI regulatory fragmentation. ***While sweeping in ambition, the EO does not impact obligations under existing state AI laws. Companies developing and distributing AI offerings should continue to comply with all existing state AI requirements***.

State officials and advocates have responded to the EO with strong criticism and early positioning for legal challenges, arguing that the EO overreaches on states' traditional police powers and consumer protection authority, and vowing to contest the EO's directive to federal agencies of creating nationwide rules that may preempt state AI laws. ***The EO is likely to be litigated, particularly its use of federal funding conditions and federal agency directives to discourage or invalidate the implementation of state AI laws***.

For the time being, states retain broad authority to enforce existing AI regulations. Further, under generally applicable consumer protection "unfair or deceptive acts and practices" and anti-competition statutes, state attorneys general and regulators may continue to pursue investigations and enforcement actions based on alleged deceptive, misleading, discriminatory, or unfair AI practices, even where those claims are framed outside AI-specific statutes. ***As a result, the EO does not limit state enforcement risk in the near term, particularly where AI deployments involve how AI products or services are marketed to consumers, how automated decisions are made, or how personal data is collected and used. State authorities can still bring enforcement actions against violators of generally applicable state statutes (implicating, e.g., consumer privacy and consumer protection)***.

International AI Regulations

Outside the United States, the EU is developing a comprehensive framework that governs AI systems, automated decisions, and the data relied on for such AI decisions through three layered regulations. The ***EU AI Act*** classifies AI systems by

risk, imposing strict requirements on “high-risk” applications (e.g., financial or employment) that cover data quality, transparency, and human oversight. General-purpose AI models must meet basic transparency standards, while the most powerful and widely deployed versions (known as “systemic models, e.g., advanced large language or multimodal models used widely across many sectors) face additional obligations for testing, safety evaluations, and incident reporting to reduce the chance of broad, society-wide harm.

Separate from the EU AI Act, ***GDPR Article 22*** restricts decisions based solely on automated processing, including profiling, that produce legal or similarly significant effects (e.g., credit approvals, hiring, or access to key services). Such automated decisions are permitted only in limited cases (e.g., with consent, or as required by contract or law) and must include safeguards such as the right to human intervention, to express a view, and to contest the decision.

The ***EU Data Act*** complements the EU AI Act and the GDPR by giving users broad rights to access and share data generated by connected products and related services. The EU Data Act requires “access-by-design,” mandates fair terms for business-to-business data sharing, and imposes cloud-switching and interoperability obligations on data-processing service providers to reduce the risk of “vendor lock-in” and make it easier to transfer and reuse data across service providers. Together, these EU measures reflect a broader global trend: countries are moving from voluntary guidelines toward enforceable, risk-based regulations designed to capture AI’s benefits while regulating opaque, high-impact, and data-intensive applications that pose the greatest risks of social harm.

Law	Effective Date	Application	Core Requirements	Enforcement and Penalties
EU AI Act	The EU AI Act takes effect in phases. It entered into force on August 1, 2024, with a transition period before most rules apply. The ban on prohibited “unacceptable risk” AI practices came into effect in early 2025, followed later that year by requirements for general-purpose AI models and governance structures. Most remaining obligations, including those for high-risk AI systems, come into force on August 2, 2026. Certain specialized and legacy systems have	Covers: ***“AI systems”*** placed on the EU market or put into service in the EU, including by providers outside the EU whose systems are used in the EU. For employment, Annex III classifies as high-risk AI used for recruitment, selection, hiring, promotion, termination, task allocation, performance monitoring, and evaluation of behavior or conduct of individuals in work-related relationships.	Core requirements under the EU AI Act include: • Classifying AI systems by risk (from minimal to unacceptable) and prohibiting certain “unacceptable risk” practices, such as manipulative social scoring or some forms of biometric surveillance. • ***Imposing obligations on “high-risk” AI systems, including a risk-management system, high-quality training data, technical documentation and logging, transparency to users, human oversight, and appropriate accuracy,***	Supervisory and market-surveillance authorities in EU Member States can investigate AI systems, order remediation or withdrawal from the market, and impose administrative fines. Depending on the type and severity of the violation, fines can reach up to €35 million or 7% of global annual turnover for certain prohibited practices, with lower tiers (for

	compliance deadlines stretching toward 2030.		***robustness, and cybersecurity.*** • Treating AI used in employment, worker management, and access to self-employment as high-risk, which triggers additional duties for providers and deployers around risk assessment, worker information, oversight, and monitoring for discriminatory impacts. • Establishing specific transparency obligations for certain AI systems (such as AI chatbots), so users are informed they are interacting with AI or viewing AI-generated content. • Introducing tailored requirements for general-purpose AI models, with baseline transparency and documentation duties and, for the most capable general-purpose AI models with systemic risk, additional testing, risk-management, and incident reporting obligations. Requiring providers and deployers of high-risk AI systems to register them in EU databases, undergo conformity assessment, affix the CE marking, carry out post-market monitoring, and take corrective actions or report serious incidents where necessary.	example, up to €15 million or 3% for breaches of the obligations applicable to high-risk or general purpose AI, and €7.5 million or 1% applying to other violations and to the provision of misleading information to authorities).
GDPR Article 22 (Automated decision-making affecting individuals)	In force since May 25, 2018, as part of the EU General Data Protection Regulation (GDPR), and directly applicable in all EU Member States. The UK has retained the GDPR in domestic law through the "UK GDPR," which largely mirrors the EU GDPR but is now interpreted and amended by UK institutions after Brexit, and sits alongside the UK's Data Protection Act of 2018.	Applies to controllers that carry out "solely automated" decision-making, including profiling, that produces legal effects or similarly significant effects on individuals, such as employment-related decisions (e.g., automated rejection of applicants, promotion/termination decisions, or salary/shift allocation driven only by algorithms).	• ***In most cases, individuals have the right not to be subject to a decision based solely on automated processing, including profiling, when those decisions have legal or similarly significant effects, unless a specific exception applies (such as contractual necessity, EU/member-state law, or explicit consent).*** • Where an exception applies, controllers must implement safeguards including giving individuals the right to obtain human intervention, to express their point of view, and to contest the decision, and controllers must also provide clear and transparent information about the logic involved, as well as the significance and anticipated consequences	Supervisory authorities enforce Article 22 using the GDPR's general enforcement powers, including investigations, corrective orders, and administrative fines. In serious cases, violations of data subject rights and core data protection principles can lead to fines of up to €20 million or 4% of a company's total worldwide annual turnover, whichever is higher.

			of the processing for the individual.	
EU Data Act	September 12, 2025 (main obligations begin applying from this date)	Covers “connected devices” placed on the EU market, including data generated by connected products and the services that support them (e.g., smart home devices, industrial machinery, and connected vehicles). This also includes both personal data (such as location or usage behavior) and non-personal data (such as sensor outputs and equipment performance metrics).	Core requirements under the EU Data Act include: • Granting users (individuals and businesses) extensive rights to access, use, and share data generated by their connected products and related services, including real-time access where technically feasible. • Designing connected products and related services so that users can easily and, in principle, freely access their data (“access-by-design”) and technically enable onward sharing with third parties on request. • Requiring data holders to share such data with users’ designated third parties on fair, reasonable, and non-discriminatory terms, and banning or invalidating unfair contract clauses that unduly restrict data access or overcharge for it. • ***Imposing data portability and switching obligations on cloud and other data-processing service providers, including limits on exit fees, migration support, and interoperability requirements to reduce vendor lock-in.*** • Allowing EU public bodies and institutions to request access to certain data in situations of exceptional need (such as public emergencies or specific public-interest tasks), subject to safeguards for trade secrets, security, and data protection. Requiring providers and deployers of high-risk AI systems to register them in EU databases, undergo conformity assessment, affix the CE marking, carry out post-market monitoring, and take corrective actions or report serious incidents where necessary.	Each EU country must designate one or more competent authority or authorities to monitor and enforce the Act, and where there are several, appoint a single “data coordinator” as the national one-stop shop and liaison for cross-border cases. Penalties under the EU Data Act are set by each EU Member State, but they must be “effective, proportionate and dissuasive,” and many countries are likely to align them with existing GDPR-style fine levels. Where a breach of the EU Data Act also involves the processing of personal data, authorities can rely on the GDPR’s fines, which allows administrative fines of up to €20 million or 4% of a company’s worldwide annual turnover, whichever is higher, for the most serious violations.

PRACTICAL GUIDANCE

Companies doing business in the EU must comply with overlapping AI regulations, including the EU AI Act, GDPR, and EU Data Act. Together, these laws impose new risk-based obligations on AI developers and deployers, and expose companies (even those that are solely U.S.-based) to a risk of very large fines for non-compliance.

As practical guidance, companies should consider the following:

- ***Map EU-exposed AI Systems***: Inventory AI and ADMT that involve EU residents or EU-sourced data and preliminarily classify them against likely EU AI Act risk tiers while flagging where the GDPR and the EU Data Act clearly apply.
- ***Establish an AI Governance Group***: Form a cross-functional team (i.e. legal, privacy, security, product, HR) to own policies, approve higher-risk deployments, and coordinate compliance under the EU AI Act, GDPR (including Article 22), and the EU Data Act.
- ***Embed Compliance-by-Design***: Build standardized risk assessments, documentation templates, logging, data-quality checks, and human-oversight gates into AI development and deployment for all high-impact systems.
- ***Enable User Data Rights***: Design or improve technical flows so users can access, port, and share device-generated data, and so the organization can handle access, portability, and cloud-switching requests without ad-hoc fixes.
- ***Update Contracts and Vendor Oversight***: Revise commercial and cloud agreements to cover AI and data risks (e.g., training-data provenance, audit/cooperation clauses, exit and interoperability terms, allocation of regulatory responsibilities).

Please refer to our prior client alert, Demystifying the EU AI Act, for additional details.

Comprehensive AI Governance

This section summarizes a growing set of state AI regulations that reflect a shift

toward comprehensive AI governance frameworks, particularly for AI systems used in high-impact contexts or to make "consequential decisions." These laws move beyond sector-specific or disclosure-only requirements and instead impose affirmative obligations related to risk assessment, governance, documentation, and oversight. While the statutes vary in scope and enforcement posture, taken together, they signal increasing regulatory expectations for technology companies that develop, deploy, or operationalize AI systems. The table below highlights key points to help companies assess near-term compliance priorities and longer-term governance strategy.

Bill	Effective Date	Application	Core Requirements	Enforcement and Penalties
California SB 53 (Transparency in Frontier AI Act)	January 1, 2026, with staggered implementation dates for covered frontier developers	Applies to "covered frontier AI developers" (e.g., developers of models with large training compute thresholds) that train or substantially modify frontier models in California (i.e. not downstream businesses users or deployers of large-scale AI systems).	Covered developers must: • Adopt and publish a "frontier AI framework" describing how they identify and manage material risks associated with AI models, including risks of misuses, systemic harms, or other such severe safety issues. Covered developers must publish within a specified period after a qualifying training run, and maintain updates after training any new frontier models. • Implement internal safety, security, and incident response measures (e.g., pre-deployment testing, red-teaming, monitoring, and controls on model access), and periodically update and review such measures with each release. • Implement processes for employee reporting and whistleblowing about safety concerns, including protections against employer retaliation.	No private right of action. Up to $1,000,000 per violation, which will be scaled to the severity of the violation (e.g., systemic failures). SB 53 is enforceable by the California AG and certain other public authorities, with the right to seek injunctive relief and civil penalties for violations.
Colorado SB 24-205 (Colorado AI Act)	June 30, 2026 (originally February 1, 2026)	Broadly applies to businesses of all sizes operating as "developers" or "deployers" of "high-risk AI systems" used to make "consequential decisions" about an individual (decisions affecting, e.g., employment, education, financial services, healthcare, housing, etc.).	Imposes a duty of care for developers and deployers to prevent algorithmic discrimination. Both developers and deployers are required to notify the Colorado AG of AI systems capable of making a discriminatory decision. • ***Developers***: provide documentation to deployers, post public website notice about AI system, conduct impact assessments. • ***Deployers***: conduct impact assessments, implement risk management policies, post public website notice about AI system, provide consumer notices and disclaimers about AI system, implement consumer opt-out rights, establish consumer appeal	No private right of action. Up to $20,000 per violation. Colorado AG has exclusive enforcement authority, and penalties can be assessed pursuant to Colorado's statutory protections against unfair and deceptive trade practices. Maintenance of impact assessments and documentation necessary to benefit from safe harbor protections.

			mechanisms for adverse consequential decisions.	
Texas HB 149 (Texas Responsible AI Governance Act)	January 1, 2026	Primarily applies to TX governmental entities and, in more limited respects, businesses that develop, deploy, or use covered "AI systems" in TX.	TRAIGA is narrower than the Colorado AI Act and does not establish a general duty of care or a broad algorithmic discrimination framework for businesses. Business requirements include: • Must provide "clear and conspicuous" notice when individuals interact with an AI system in specified contexts, particularly where AI simulates human interaction or materially influences outcomes. • Prohibition on certain uses of AI, including biometric identification, social scoring, and other expressly restricted AI applications.	No private right of action. Civil penalties of up to $10,000 per violation, with right to seek injunctive relief. Texas AG has exclusive enforcement authority, and penalties can be assessed pursuant to Texas's unfair and deceptive trade practices statutes.
Montana SB 212 (Right to Compute Act)	October 1, 2025	Applies broadly to providers of digital services, software, or computing resources that operate in or affect users in Montana. Note that the Act is not an AI-specific law, but it establishes a statutory "right to compute" for users and businesses.	The Act prohibits covered providers from restricting, degrading, or interfering with lawful computing activity, including the ability to: run lawful software or algorithms of the user's choosing, access computing resources necessary to perform lawful computational tasks, or use computing power for purposes such as data analysis, cryptography, AI model training, or other lawful compute-intensive activities. Providers may impose restrictions only where necessary to: comply with federal or state law; preserve system integrity, security, or reliability; prevent demonstrable harm, fraud, or abuse (provided that such restrictions are narrowly tailored).	No private right of action. Civil penalties may be assessed for violations, including injunctive relief and monetary penalties under Montana's consumer protection statutes. Montana AG has exclusive enforcement authority.
New York S6953-B (NY RAISE Act)	January 1, 2027	Does not apply to most businesses. The Act narrowly targets catastrophic harm caused by developers of "frontier models." Application is limited to: (1) AI developers with more than $500M in annual revenue, and (2) companies that develop or operate frontier AI models in NY.	Imposes affirmative obligations on developers and deployers to identify, assess, and mitigate risks of algorithmic discrimination and other foreseeable harms associated with high-risk AI systems. Key frontier model developer obligations include: • Create and follow written safety protocols, conduct AI impact assessments on risk of "critical harm" to persons or property, and implement appropriate safeguards. • Report incidents to the NY AI oversight office within 72 hours of determining that an incident has occurred. The Act also creates a new AI oversight office for registration, assessment of oversight fees, creation of new regulations/guidance, and publication of annual reports on AI safety risks.	No private right of action. New York AG has enforcement authority, and can levy penalties of up to $1M for the first violation, or up to $3M for subsequent violations.

PRACTICAL GUIDANCE

Note that, while many of these laws include applicability thresholds that early-stage or mid-market technology companies may not meet directly, startups may still face indirect compliance and contracting risks when using AI tools or services from vendors that are subject to these regulations. In practice, this dynamic is increasingly reflected in AI-specific addenda and contractual terms with expansive representations, warranties, or compliance assurances relating to a company's use of third-party AI tools. As a matter of best practice, companies should exercise caution when asked to make guarantees beyond their visibility or control as a matter of best practice.

Your Gunderson team can provide practical compliance resources, including impact assessment templates and internal governance toolkits, to help your company evaluate applicability, operationalize requirements, and plan for scalable AI compliance. Additionally, please refer to AI in the Workplace: Legal Challenges and Best Practices and AI Regulatory Landscape Under the New Trump Administration, or subscribe to updates on our AI Resources portal, for prior and future client alerts and webinars covering such comprehensive AI governance regulations.

Consumer AI Interactions

This section summarizes a new wave of state laws targeting consumer-facing AI interactions, especially chatbots and algorithmic pricing systems. Together, these laws reflect a growing consensus that AI systems which simulate human-like interaction or tailor prices using personal data must meet enhanced transparency and safety expectations. While their specific triggers and remedies differ, these laws generally focus on clear disclosures that users are dealing with AI, restrictions and protocols around high-risk uses, and enforcement through state consumer protection authorities or opening the door to private litigation.

Bill	Effective Date	Application	Core Requirements	Enforcement and Penalties
California SB 243 (AI Companion Chatbot Safety)	January 1, 2026	Targets "companion chatbots" (e.g., AI systems that provide adaptive, humanlike responses and capable of maintaining ongoing, relationship-style user interactions).	Establishes requirements on operators to provide disclosures and notices, safety protocols, protections for minors and against harmful content, and to implement monitoring/reporting governance functions. These include:	Creates a private right of action. Any person who suffers an "injury in fact" from a violation may bring a civil action against an operator. Available

		Applies to "operators" that make companion chatbots available to users in California, with special emphasis on use cases involving minor and emotional/mental wellness.	• ***Disclosures and Notices:*** Must clearly and conspicuously disclose that users are interacting with AI at the beginning of each interaction, with reminders every 3 hours and suitability warnings when users are minors. • ***Safety Protocols:*** Protocols to detect self-harm and suicide, including referrals to crisis prevention providers (e.g., suicide hotlines or crisis text lines). • ***Content Restrictions for Minors:*** Reasonable measures to prevent sexually explicit conduct, and to avoid coercive engagement-maximizing tactics. • ***Monitoring and Reporting:*** Beginning July 1, 2027, operators must maintain records and report crisis-related interactions (including data on crisis incidents	remedies include injunctive relief, damages equal to the greater of actual damages or a statutory minimum of $1,000 per violation, plus reasonable attorneys' fees and costs. California AG also has enforcement authority under California's consumer protection and unfair competition laws.
New York S-3008C, Part U (AI Companions)	November 5, 2025	Targets "AI companions" (e.g., AI systems that simulate a sustained humanlike relationship, capable of maintaining and engaging a simulated conversation on personal wellbeing). Applies to any "operator" that operates or provides an AI companion to users in New York.	Establishes requirements on operators, including: • ***Disclosures and Notices:*** Must clearly and conspicuously disclose that users are communicating with an AI chatbot (i.e. not a human) at the beginning of each interaction, and at least once every 3 hours in ongoing interactions. • ***Safety Protocols:*** Operators must include an AI companion protocol that takes reasonable efforts to detect and address users' expressions of suicide ideation or self-harm, including notifying and referring users to crisis prevention providers (e.g., suicide hotlines or crisis text lines).	No private right of action. Up to $15,000 per day for violations of notification and safety protocol requirements, and directs collected penalties into a dedicated suicide prevention fund. New York AG has enforcement authority, including seeking injunctive relief and civil penalties.
New York S-3008C, Part X (Algorithmic Pricing)	November 10, 2025	Applies to entities domiciled or doing business in New York that determine prices for goods or services using "personalized algorithmic pricing" (e.g., dynamic pricing derived from or set by an algorithm using consumer personal data). Limited exemptions for financial institutions, and subscription-based pricing where the algorithmic price is lower than the consumer's existing subscription price.	Covered entities must provide a clear and conspicuous disclosure alongside any personalized algorithmic price, using the exact wording: "THIS PRICE WAS SET BY AN ALGORITHM USING YOUR PERSONAL DATA." Use of certain "protected class data" (e.g., data linked to protected characteristics) in pricing is restricted where it would result in denial or withholding of accommodations, advantages, or privileges, or in different prices for such groups; targeted marketing using protected-class data is still permitted if it does not cause groups to miss benefits like discounts.	No private right of action. Up to $1,000 per violation. New York AG has enforcement authority, including right to seek injunctive relief and civil penalties. New York AG also has the right to issue cease-and-desist notices, and provide opportunities for businesses to cure violations.
Maine LD 1727	September 23, 2025	Applies to any "person" using an "AI chatbot" (e.g.,	A person may not use an AI chatbot (or other covered technology) to	No private right of action.

(AI Chatbot Consumer Disclosures)		app or program that simulates human conversation and interaction through text or voice) to engage in trade or commerce with a consumer in Maine.	engage in trade or commerce with a consumer in a way that may mislead or deceive a reasonable consumer into believing they are interacting with a human, unless the consumer is notified in a clear and conspicuous manner that they are not engaging with a human being.	Violations are enforced under the Maine Unfair Trade Practices Act, and exclusively enforced by the Maine AG. Remedies and penalties are provided pursuant to Maine's Unfair Trade Practices Act.
Utah SB 226 (AI Consumer Protection Amendments)	May 7, 2025	Targets AI used in consumer transactions and certain regulated occupation services. Introduces a defined concept of "high-risk AI interaction" (e.g., collection of sensitive personal data; or provision of recommendations that could be relied upon for significant decisions, such as financial, legal, medical, or mental health advice).	Establishes disclosure-driven consumer protection obligations, including: • ***Suppliers:*** If a consumer interacts with AI in connection with a transaction and asks/prompts the supplier about whether AI is being used, the supplier must provide a "clear and conspicuous" notice that the consumer is interacting with AI (i.e. not a human). • ***Regulated Occupation Services:*** Service providers in regulated occupations must prominently disclose when the consumer is interacting with AI if the use constitutes a high-risk interaction. The disclosure timing is specified (i.e. verbal at start of verbal interaction; in writing before start of written interaction) and throughout the applicable regulated AI interaction.	No private right of action. Up to $5,000 per violation. Utah AG has enforcement authority, and penalties can be assessed pursuant to Utah's consumer protection laws. Clear and conspicuous disclosure at the outset and through an AI interaction is necessary to benefit from safe harbor protections.

PRACTICAL GUIDANCE

Companies that provide or operate consumer-facing AI interactions should practice product, governance, and contract hygiene in a manner that scales with regulatory risk. For example:

- ***AI Inventory:*** Understand and document where products or services involve AI functionalities (e.g., chatbots, personalized/dynamic pricing, etc.) and identify which flows are consumer-facing vs. internal. Treat any product feature that simulates human conversation, provides emotional support, or adjusts pricing at the individual level as in-scope for heightened internal review.
- ***Disclosure and Notification:*** Build clear, conspicuous, and persistent disclosures into consumer product flows, especially where users could reasonably think they are interacting with a human or where user-specific prices are set using personal data. Use plain English and test disclosures in context (e.g., multi-session conversations, mobile, minor access, etc.) to avoid any regulatory scrutiny of "dark pattern" behavior.
- ***Implement High-Risk Safety Protocols:*** For use cases involving personal or mental health, handling minor access to content, or affecting significant financial/legal decisions, implement frameworks to address legal risk. These may include content filters, escalation logic, crisis detection and messaging, and rate-limiting or cooling-off mechanisms for prolonged AI use. Document and maintain records, and implement periodic testing of existing protocols.
- ***Review Vendor and Customer Contracts:*** Ensure that contracts with providers of third-party AI models include transparency obligations, safety filters, procedures for handling disclosures to state authorities, and commercially reasonable indemnification obligations and validation processes. Customer terms must align with provided company controls and compliance requirements.

Additionally, please refer to California SB 243: New Compliance Requirements for Operators of AI Companion Chatbots, or subscribe to updates on our AI Resources portal, for prior and future client alerts and webinars covering use-specific consumer AI regulations.

AI Content Transparency

This section highlights emerging AI content transparency regimes that focus on how AI-generated content is created, labeled, and presented to consumers. Together, these regulations signal a shift from general AI governance toward targeted provenance, disclosure, and warning-label obligations for AI-generated media, synthetic performers, and addictive content experiences. Although they vary in scope and enforcement posture, these laws collectively raise expectations that developers, platforms, and advertisers will surface when AI is used to generate or alter content, and provide clear warnings (particularly to minor users) backed by public enforcement and, in some cases, significant per-violation penalties.

Bill	Effective Date	Application	Core Requirements	Enforcement and Penalties
California AB 2013 (AI Training Data Transparency)	January 1, 2026	Applies to any person or entity that designs, codes, produces, or substantially modifies a generative AI system made available (i.e. free or paid) for public use by California residents. Covers generative AI systems first released or updated on or after January 1, 2022, that can generate synthetic content such as text, images, audio, or video.	Developers must post on a publicly accessible website a "high-level summary" of training data for each covered generative system, including: • Data sources and ownership; • Types and volume of data; • Collection and processing methods; • Copyright, trademark, patent, or public-domain status and license details; • Whether datasets include personal or aggregate consumer information (as such terms are defined under the California Consumer Privacy Act); • Collection timeframes and first-use dates; and • Whether any synthetic (i.e. AI generated) data was used in training or development. The summary must be updated each time the developer substantially modifies the AI system, such as by training or expanding datasets in a manner that materially alters capabilities.	No private right of action. Silent on enforcement of AB 2013 obligations, and does not designate a lead California enforcement agency. Likely enforcement under California's Unfair Competition Law or related consumer protection authorities, especially if a developer makes false or misleading statements about training data.
California AB 853, as amended (California AI Transparency Act (CAITA))	August 2, 2026; obligations for AI hosting platforms begin January 1, 2027, and capture-device manufacturers on or after January 1, 2028	"Covered providers" include: (1) providers that create, code, or produce generative AI systems with more than 1,000,000 monthly users in California; (2) social media or content platforms distributing AI content to California users; (3) AI hosting platforms that make	• ***AI Developers:*** Must (1) provide a free AI-detection tool capable of identifying latent disclosures in AI-generated or -altered multimedia content; (2) embed latent manifest disclosures in AI-generated content to convey provenance information (e.g., provider/system name,	No private right of action. $5,000 per violation, with each day of non-compliance treated as a separate violation. California AG and local regulatory agencies have enforcement authority, with

		source code or model weights available for download by California residents; and (4) producers of devices capable of recording sold in California.	creation/alteration timestamps, unique identifiers, etc.); and (3) contractually require third-party licensees to maintain latent disclosures in content. • ***AI Hosting Platforms (after January 1, 2027):*** Must not knowingly make an AI system that lacks manifest latent disclosures required under CAITA (i.e. only host systems that support AI provenance markers). • ***Capture-Device Manufacturers (after January 1, 2028):*** Must provide user option to add a latent disclosure in content captured by the device and, by default, embed disclosures in captured content to support authenticity verification of human-captured media.	enforcement expected to proceed under California's Unfair Competition Law and related consumer protection regulations.
New York SB-8420A (Synthetic Performers in Advertising)	June 9, 2026	Applies to any person who "produces or creates an advertisement" for a commercial purpose, in any medium, with actual knowledge that a "synthetic performer" appears. Statute excludes certain creative and editorial uses, including expressive works where the synthetic performer's use is "consistent with its use in the underlying work" (e.g., film, TV, video games, etc.).	Any covered visual or mixed-media commercial advertisement that includes a "synthetic performer" (e.g., digitally-created AI asset that is intended to create the impression of a human performer) must conspicuously disclose within the advertisement itself that a synthetic performer is being used. The law does not prescribe exact wording or specific disclosure format, but the disclosure must be clear and noticeable and adapted to each medium where the advertisement is run (e.g., on-screen text for visual ads, audible statements for audio-only ads, and clear labeling for digital or social placements, etc.).	No private right of action. $1,000 for the first violation and $5,000 for each subsequent violation of the disclosure requirement. Enforcement is expected to proceed through New York state authorities under the NY General Business Law, with advertisers and agencies primarily responsible for meeting statutory compliance requirements.
New York S4505/A5346 (Social Media Labeling)	TBD; operative in 2026 after NY Commissioner of Mental Health's publication of labeling standards.	Applies to "addictive social media platforms" that provide personalized feeds, autoplay, infinite scroll, and/or push notifications as a significant part of the service. Protects covered minors and covers platforms making such features available to users in New York.	While this is not strictly an AI law, the breadth of the regulation could apply to companies using AI algorithms to personalize social media feeds. The law's protections are keyed to "young users," defined as minors "reasonably known" to be under 18, including accounts self-declared as minors or identified via age-assurance tools, with specific emphasis on enhanced warnings and usage triggers for users under 18. Covered platforms must display a clear, conspicuous warning label that alerts young users to the potential mental health risks associated with the platform's addictive features, using language prescribed by the Commissioner of Mental Health (such language remains forthcoming), upon signup and periodically thereafter based on continued use. This includes showing a warning after 3 cumulative hours of active use in a day, and at least once per additional hour.	No private right of action. $5,000 per violation (i.e. failure to present a required warning to a covered young user as prescribed by NY regulatory authorities). New York AG has exclusive enforcement authority, and may seek injunctive relief.

PRACTICAL GUIDANCE

Companies providing generative AI tools should treat provenance and labeling as a core product requirement (i.e. design persistent, understandable indicators when content is AI-generated, synthetic performers are used, or minors are exposed to “addictive” features), which may include:

- Building or adopting industry-standard technical tooling for manifest and latent disclosures;
- Setting clear internal/external policies for when labels, warnings, and training-data summaries are required;
- Training marketing, product, and engineering teams on applicable thresholds;
- Periodically sampling ads, product surfaces, and published AI summaries to confirm that required notices are present and current, documenting changes as laws and guidance evolve.

Further, companies should be ready to adapt to forthcoming prescriptive guidance from regulators as well as potential outcomes from court decisions. For example, the New York Commission of Mental Health is expected to release required warning label language for minors using social media features. Also, on December 29, 2025, xAI (formerly Twitter) filed suit against the California AG to enjoin enforcement of California AB 2013, alleging that the law’s training disclosure requirements (1) are an “unconstitutional taking” under the Fifth Amendment that forces xAI to disclose valuable trade secrets without fair compensation, and (2) compel speech in violation of the First Amendment. ***Companies will need to continue monitoring for legal challenges to implementation of these state laws, as well as publication of content disclosure requirements by relevant state authorities.***

Employment and Automated Decision-Making Tools

Automated employment decision tools are an early focal point of emerging AI regulation, as legislators and regulators move from high-level principles to prescriptive rules governing how AI may be used in hiring and workforce management. These laws and proposals generally treat automated employment

decision tools ("AEDTs") and automated decision-making tools ("ADMTs") as distinct, high-risk systems and layer new duties, such as independent bias audits, structured risk assessments, and detailed applicant notices, on top of existing anti-discrimination and privacy frameworks. While the specific triggers, timelines, and enforcement mechanisms vary by jurisdiction, the common theme is an expectation that employers and vendors can demonstrate that their tools are explainable, monitored for disparate impact, and subject to meaningful human oversight. The AEDT developments summarized below are intended to help organizations prioritize near-term compliance steps while building longer-term AI governance programs that can withstand evolving scrutiny from regulators, courts, and stakeholders.

Bill	Effective Date	Application	Core Requirements	Enforcement and Penalties
NYC Local Law 144 (AEDTs)	July 5, 2023	Applies to employers and employment agencies that use AEDTs to substantially assist or replace discretionary decision-making for hiring or promotion decisions impacting candidates or employees in New York City.	Covered entities must: • Conduct an independent bias audit of each AEDT within one year before use, and repeat at least annually. • Publish a summary of the most recent bias audit (including key metrics and date of distribution) on the employer's or agency's website before using the AEDT. • Provide candidates and employees who reside in NYC with notice at least 10 business days before the AEDT is used, describing that an AEDT will be used, what qualifications/characteristics it assesses, types and sources of data, and data retention policies, and offering the opportunity to request an alternative process or accommodation.	No private right of action. The NYC Department of Consumer and Worker Protection (DCWP) enforces the law, including audit, notice, and publication obligations. DCWP may impose civil penalties of up to $500 for a first violation and $500–$1,500 for each subsequent violation, with each day of unlawful AEDT use and each failure to provide required notice treated as a separate violation. Enforcement challenges, exceptions, and uncertainty have resulted in minimal compliance to date.
California CPPA/CPRA (ADMTs)	January 1, 2027 (with some obligations beginning 1.1.26)	Regulations apply to businesses subject to the CCPA/CPRA that use ADMTs to make or materially influence "significant decisions" about consumers, including in employment, credit, housing, education, and similar high-impact contexts. ADMT is defined broadly as technology that processes personal information to execute or substantially facilitate decisions, including profiling and AI systems used in employment decision-making.	Covered entities must: • Conduct and document a risk assessment of any high-risk ADMT before first use for tools deployed on or after January 1, 2026, recognizing that ADMT-specific obligations take full effect in 2027 and that formal risk-assessment submissions begin in 2028. • General compliance by January 1, 2026; full ADMT compliance by January 1, 2027; attestations due April 1, 2028. • Provide pre-use notices when ADMT is used for significant decisions, explaining the purpose of use, the logic involved in the ADMT, and material factors	No private right of action. The California Privacy Protection Agency and the California AG may enforce violations of the CCPA/CPRA and its ADMT regulations through administrative enforcement and civil actions. Businesses face statutory penalties of up to $2,500 per violation or $7,500 per intentional violation or violations involving minors, along with injunctive relief and mandated remedial measures.

			considered, and offering a right to opt out where required. • Offer consumers the ability to access meaningful information about the ADMT's functioning and, in specified contexts, to contest or seek human review of ADMT-driven decisions. • Conduct written risk assessments for ADMT used for significant decisions or sensitive profiling at least annually, documenting purposes, data categories, risks to consumers, and safeguards; and submit these assessments to the CPPA.	
California Civil Rights Regulations (ADS in employment context)	October 1, 2025	Regulations apply to employers with 5+ employees using automated tools, algorithms, or AI systems to make or substantially assist employment-related decisions under California's civil rights laws. The rules are framed within anti-discrimination and fair employment statutes, focusing on tools that could result in disparate treatment or impact in hiring, promotion, or other employment actions.	Covered entities must: • Ensure that automated decision systems used in employment do not directly or indirectly discriminate based on protected characteristics, aligning ADS use with existing anti-discrimination obligations. • Implement governance measures such as testing, validation, and ongoing monitoring of ADS for discriminatory impact, and adjust or discontinue tools that cause unlawful disparate impact. • Provide appropriate pre-use notices and accommodations where ADS interacts with applicants or employees, consistent with broader civil rights and fair employment rules.	Individuals can bring claims under California civil rights statutes. The California Civil Rights Department (and Civil Rights Council through rulemaking) enforces these requirements using the same mechanisms available under state civil rights law, including investigations, administrative complaints, and civil actions. Remedies may include injunctive relief, hiring or reinstatement orders, back pay and damages, and civil penalties for patterns or practices of discrimination.
Illinois AIVIA (AI Video Interviews)	January 1, 2020	Applies to employers that ask applicants to record video interviews and use AI to analyze those videos when considering applicants for positions based in Illinois. The law is limited to video-based screening and is technology-specific, targeting AI systems used to assess video interviews for employment purposes.	Covered entities must: • Inform applicants before their interview that AI will be used to analyze their video and explain in general terms how the AI works and what characteristics it evaluates. • Obtain the applicant's consent before using AI to evaluate the video; if consent is not provided, the employer may not use AI for that applicant's interview. • Restrict sharing of interview videos to persons whose expertise or technology is necessary to evaluate the applicant and delete the video (and require others to delete copies) within 30 days of an applicant's request.	No private right of action. The Act itself does not specify penalties, remedies, or a dedicated enforcement mechanism, leaving open questions about whether violations can be pursued via implied private rights, agency enforcement, or through other Illinois statutes.
Maryland HB 1202 (Use of facial recognition)	October 1, 2020	Prohibits employers from using facial recognition services to create a facial template during a job applicant's interview unless the applicant signs a written consent waiver.	Covered entities: HB 1202 is limited to facial-recognition during interviews; it does not establish a comprehensive AEDT framework covering other types of automated hiring tools.	No explicit private right of action. Any private claim would likely have to be framed under other theories (e.g., common-law or existing discrimination statutes), and

				there is no clear case law yet on such claims.
Maryland HB 1255 (Employment-focused AEDTs	TBD	Would restrict employers' use of "automated employment decision tools" for certain employment actions, require notice to applicants, and condition use on an impact assessment showing the tool does not result in unlawful discrimination or disparate impact. A related Senate bill (SB 957) would similarly prohibit, with limited exceptions, the use of automated employment decision tools to make specified employment decisions; as of the latest updates these bills had not been enacted and were still moving (or stalled) in the legislative process.	• Require employers that use automated employment decision tools to give notice to applicants or employees that such tools will be used in connection with specified employment decisions. • Require an "impact assessment" or similar evaluation to determine whether use of the tool results in discriminatory or disparate impact, with related record-keeping obligations that effectively document the tool's fairness and reliability. • Direct state agencies to adopt implementing regulations, which are expected to include standards for bias assessment, documentation, and periodic monitoring of covered tools, thereby formalizing ongoing oversight expectations for employment-related ADMT	No private right of action beyond existing Maryland civil rights and employment law. Enforcement would generally align with existing Maryland labor and anti-discrimination enforcement structures, allowing regulators to investigate and enforce compliance with transparency and fairness requirements. Penalties would likely include civil fines and injunctive relief, although specific dollar ranges and private rights of action depend on the final enacted text or subsequent amendments.

PRACTICAL GUIDANCE

Automated employment decision tool laws are already influencing how vendors design hiring products and how employers structure procurement and oversight, even where a particular statute does not yet apply directly to a company. Many regimes place primary obligations on the "user" or "deploying" employer, but vendors are increasingly asked to support compliance by furnishing bias audits, documentation, and technical controls, which can create back-pressure on startups that supply or embed AEDTs in their products. As these requirements expand, companies should map where automated tools meaningfully influence employment decisions, tighten internal approval and review processes for new tools, and build repeatable practices for evaluating bias, documenting testing, and responding to regulator or customer inquiries. From a contracting standpoint, it is prudent to resist open-ended representations about fairness or legal compliance, instead tying obligations to documented controls and agreed-upon testing protocols, and ensuring that indemnities and limitation-of-liability clauses reflect the heightened risk profile of employment-related AI systems.

Additionally, please refer to AI in the Workplace: Legal Challenges and Best Practices, Quarterly Employment Law Update – Summer 2025 (for brief discussion of California's Civil Rights Regulations and ADS) and Legislating the Future of AI in Employment: NYC's Law on Automated Decision Tools & Other Important Developments. Your Gunderson team can assist with testing and auditing your AEDTs, and with creating strategies for using AI in the workplace. Please also subscribe to updates on our AI Resources portal for prior and future client alerts and webinars covering AI and employment law.

Anti-Discrimination and Civil Rights

Across jurisdictions, lawmakers and regulators are increasingly framing AI issues as extensions of long-standing anti-discrimination and civil-rights principles in the workplace. They are making clear that employers remain responsible for biased or exclusionary outcomes even when those outcomes are produced by algorithms or third-party tools, not explicit human decisions. Guidance and emerging bills emphasize familiar concepts, such as disparate impact, reasonable accommodation, and vendor oversight, but apply them to automated screening, scoring, and monitoring systems. The result is a growing expectation that employers will

proactively test and monitor AI systems for discriminatory effects, document their findings, and adjust or abandon tools that create unacceptable legal or equity risks. The table below highlights several recent and ongoing examples of this guidance and legislation.

Regulation/Guidance	Effective Date	Application	Core Requirements	Enforcement and Penalties
New Jersey DCR Guidance (AI discrimination across sectors)	January 9, 2025	Applies to all entities already covered by the New Jersey Law Against Discrimination (LAD), including employers, housing providers, lenders, schools, and places of public accommodation. It does not create new statutory duties; instead, it clarifies that "algorithmic discrimination" from automated decision-making tools (including AI, machine learning, and statistical models) is treated the same as any other discriminatory practice under the LAD.	The Guidance: • Emphasizes that covered entities remain fully responsible for discrimination caused by automated tools, even when developed or operated by third-party vendors, and cannot contract away LAD obligations. • Highlights that AI tools can create disparate treatment, disparate impact, and failure-to-accommodate violations, and urges entities to take steps such as testing for bias, monitoring outcomes, ensuring accessibility/ accommodations, and maintaining oversight of AI vendors.	Existing private right of action under LAD. The Guidance itself is non-binding but signals enforcement priorities; alleged "algorithmic discrimination" is investigated and prosecuted using existing LAD mechanisms by the New Jersey AG and Division on Civil Rights. Violations can result in the full range of LAD remedies, including injunctive relief (changing or abandoning tools), monetary damages, civil penalties, and attorneys' fees, with both agency and private-plaintiff actions available.
New Jersey A. 3854 (pending; AEDT discrimination)	TBD	Pending bill that would regulate the use of AEDTs in employment decisions to minimize discrimination. As of the most recent texts and commentary, A. 3854 has been introduced and revised in committee but has not been enacted, so it has no effective date yet.	Core elements of the bill: • Would cover employers and employment agencies that use AEDTs to substantially assist or replace discretionary decision-making in employment (e.g., hiring or promotion) and entities that sell AEDTs in New Jersey. • Requires vendors to perform annual bias audits of AEDTs for compliance with anti-discrimination laws and provide audit results to purchasers at no extra cost; vendors must also disclose that the tools are subject to these audits. • Requires employers using covered AI tools to obtain specified demographic data on applicants, provide required notices, and submit data to the Department of Labor and Workforce Development for state-run bias reviews.	Authorizes state labor and civil-rights authorities to enforce the bill through investigations and administrative actions, with obligations layered on top of existing LAD duties. Drafts contemplate civil penalties for non-compliance (e.g., failure to conduct audits or report data), but exact amounts and any private right of action would depend on the final enacted version; at present these provisions are only proposed and not yet in force.
New Jersey A. 3911 (pending; AI and civil rights)	TBD	Pending bill that would regulate the use of artificial-intelligence-enabled video interviews in hiring, supplementing New Jersey's	Core features of the bill: • Applies to employers that use AI-enabled video interviewing tools when	Draft language would authorize monetary penalties ranging from $500 for a first offense to $500–$1,500 for

		employment and civil-rights framework. The bill was introduced in February 2024 and remains in the legislative process without enactment.	evaluating job applicants for positions in New Jersey. • Would require employers to notify applicants that AI will analyze their video interview, explain in general terms how the AI works and what characteristics it evaluates, and obtain the applicant's written consent before using the technology. • Limits sharing of interview videos to persons whose expertise or technology is needed to evaluate candidates and may require deletion of videos within defined timeframes or upon request, similar in concept to Illinois' AI Video Interview Act.	each subsequent offense for violations of notice, consent, and use restrictions. Enforcement would occur through state labor and/or civil-rights authorities; the bill does not yet clearly establish an additional private cause of action beyond existing LAD remedies, and these enforcement details remain subject to change pending legislative action.

PRACTICAL GUIDANCE

Companies using AI in employment should treat civil-rights compliance as a primary design requirement, not a clean-up exercise. In practice, this means mapping where automated tools influence high-stakes decisions, testing for disparate impact and accessibility, and documenting both results and remediation. Employers should assume they remain responsible for vendor tools and reflect that in contracts and governance, for example by securing audit and data-access rights and avoiding broad, unverifiable "fairness" assurances. Your Gunderson team can assist with building and documenting AI fairness programs and negotiating vendor arrangements, and you can stay current on developments through our AI Resources portal.

How can GD help?

Gunderson Dettmer is a trusted partner of startups and investors and can help your company navigate evolving AI regulatory requirements, evaluate risk exposure, prioritize compliance efforts, and plan for upcoming obligations. If you have any questions regarding this client alert, or your company's compliance position with respect to these current and forthcoming AI regulations, please reach out to your Gunderson Dettmer attorney or contact any member in our Strategic Transactions & Licensing, Employment & Labor, or Privacy Groups.

Related Services

AI & Machine Learning

Data Privacy

Employment & Labor

Strategic Transactions & Licensing

Featured Insights

CLIENT NEWS

Addi Announces $85 Million Series D

CLIENT NEWS

Gunderson Dettmer Represents Ona in Acquisition by OpenAI

CLIENT NEWS

Gunderson Dettmer Represents WideField Security in Agreement to be Acquired by Cisco

FIRM NEWS

Gunderson Dettmer Welcomes Kate Cusick as Chief Communications Officer

INSIGHTS

Client Insight: Quarterly Employment Law Update – Spring 2026

CLIENT NEWS

Proception Raises $11 Million Seed Financing

CLIENT NEWS

Nebulock Raises $25 Million Series A

CLIENT NEWS

Caplight Raises $16 Million Series A

CLIENT NEWS

Gunderson Dettmer Represents Modular AI in Acquisition by Qualcomm

CLIENT NEWS

Noro-Moseley Partners Leads EAIGLE’s Growth Funding Round

CLIENT NEWS

Osanni Bio Closes $190 Million Series B

CLIENT NEWS

Gunderson Dettmer Represents Astral in Acquisition by OpenAI