Cooley

State AI Laws – Where Are They Now?

April 24, 2026

The US artificial intelligence (AI) regulatory landscape is at an inflection point. With accelerating pace starting in the 2020s, hundreds of proposed state measures signaled a fast-developing state-level regulatory AI landscape. However, as compliance deadlines in 2026 approached, many of these state AI laws that initially passed just a few years ago have undergone significant changes or delays since their passage. At the same time, federal action is potentially threatening to reshape or constrain state-level initiatives. In this alert, we check in on the current status of some of the major state AI laws.

As we discussed on March 25, the White House recently released its National Policy Framework for Artificial Intelligence, urging Congress to enact sweeping AI legislation to preempt certain state AI laws, with a focus on state laws that risk stifling innovation and avoiding "undue burdens." States like California are also leveraging executive action. For example, on March 30, 2026, California Gov. Gavin Newsom issued Executive Order N-5-26, directing state agencies to draft recommendations for AI safety requirements – including related to illegal content, bias, and civil rights and free speech – for companies doing business with state agencies. In parallel, other states are reconsidering or delaying their AI laws. Below we outline the key AI laws where companies should watch for potential changes over the coming months.

Colorado: SB 205

In May 2024, Colorado SB 205 created one of the first comprehensive state AI regimes, regulating "high-risk artificial intelligence systems" used in "consequential decisions." The law imposes broad obligations on developers and deployers related to risk management, impact assessments, consumer disclosures and reporting to the Colorado attorney general.

Since its enactment, SB 205 has been subject to significant debate and criticism, particularly from the tech industry, with concerns raised over its scope and feasibility. These concerns prompted a special legislative session in August 2025 that led to the postponement of the initial enforcement date, from February 1, 2026, to June 30, 2026.

Now, with the delayed effective date, Colorado is considering a more substantive revision. The March 2026 working group draft would repeal and reenact the newly focused law on automated decision-making technology (ADMT) and reset the effective date to January 1, 2027. The group is led by the Colorado governor's office and is composed of legislators, industry representatives, consumers and school district representatives, among others. The group was tasked with evaluating whether the original framework was workable in practice, with the ultimate goal of protecting consumers.

Key proposed changes by the working group include:

1. Replacing "high-risk AI" with "covered ADMT" that must "materially influence" a consequential decision, excluding incidental or low-stakes uses.[1]
2. Clarifying and narrowing what constitutes a "consequential decision" – specifically, limiting "consequential decisions" to high-impact decisions affecting access to education, employment, housing, financial services, insurance, healthcare or government services, where the outcome materially influences eligibility, access or opportunity.
3. Carving out routine business processes, marketing and other low-risk uses (e.g., advertising and marketing tools, recommendation and search systems, content moderation, and summarization and presentation assistance).
4. Substantially scaling back governance obligations for both developers and deployers, including eliminating requirements to implement formal risk-management programs, impact assessments, annual reviews and Colorado attorney general incident reporting. Instead, it shifts to a more targeted framework that still requires developers and deployers to maintain records and documentation regarding covered ADMT, provide consumer-facing disclosures, and implement processes for requests to correct inaccurate

information and seek human review or reconsideration of certain decisions, where commercially reasonable.
5. Replacing the pre-decision notice framework with a point-of-interaction requirement (meaning at the specific moment a user engages with the system) that may be satisfied via a prominent public posting, and adding a separate post-adverse disclosure that explains the decision, the role of ADMT and available recourse options.
6. Retaining Colorado attorney general-only enforcement but adding a 90-day notice-and-cure period and clarifying developer versus deployer liability.
7. Removing the stand-alone affirmative duty to "avoid algorithmic discrimination" that appeared as an explicit, independent requirement in the original SB 205.

These proposed, significant changes signal state action moving away from a broad "high-risk AI" framework toward a narrower, decision-focused model. With the effective date of June 30, 2026, approaching, it remains unclear whether and when the proposed amendments will be enacted. As such, companies should continue preparing for compliance under the current framework, while maintaining a watchful eye on legislative developments that could reshape obligations in the near term.

California: AB 2013, SB 942 and AB 853

California enacted 18 AI-related laws across 2023 and 2024, some of which we discussed at the time. Many of these laws impose transparency, disclosure and governance requirements on AI systems and digital services.

Given many compliance effective dates now start in 2026 and beyond, these laws have not yet seen enforcement activity or further interpretive guidance to aid in compliance; however, that may change as the year progresses.

Key California AI laws that have recently come into effect or been amended include AB 2013, SB 942 and AB 853:

- AB 2013 (training data transparency) requires developers to disclose information regarding training datasets. It was enacted on September 28, 2024, and became effective on January 1, 2026, with limited implementation guidance beyond the law's original provisions. The law requires a "high-level summary of the datasets used in the development of the generative artificial intelligence system or service," and identifies certain information for inclusion in said summary, as well as certain security-related exceptions to the disclosure requirement. Industry stakeholders have raised concerns regarding feasibility and scope, and clear patterns around the form of compliance (such as the format or level of detail for the summary) have not yet emerged. In addition, the lack of guidance or action from the California attorney general has contributed to some uncertainty on the regulatory compliance obligations.
- SB 942 (AI disclosure requirements), enacted September 19, 2024, requires providers of generative AI image, video and audio tools with more than one million monthly visitors or users to provide an AI detection tool, "manifest" disclosures (a watermarking option) and "latent" disclosures, enabling individuals to detect whether content was generated by the provider's tool. Its effective date was delayed from January 1 to August 2, 2026, via AB 853, which also added new obligations on large online platforms with an operative date of January 1, 2027, and capture device manufacturers with an operative date of January 1, 2028.
- AB 853 (California AI Transparency Act) introduces the phased implementation for SB 942 discussed above and also expands SB 942, including to add requirements applicable to large online platforms (public-facing social media platforms) and capture device manufacturers (persons producing capture devices for sale in California). These obligations include ensuring that content is appropriately labeled or identifiable as AI-generated, as well as implementing mechanisms to enable detection of such content.

Utah: SB 149

Utah's Artificial Intelligence Policy Act (SB 149), enacted on March 13, 2024, and effective on May 1, 2024, is widely viewed, together with other laws discussed in this article, as one of the first state AI governance frameworks. Rather than creating a stand-alone regulatory regime, the law primarily extends existing consumer protection principles to AI by making companies liable where AI-driven conduct would otherwise violate deceptive or unfair practices laws. The law also introduced targeted disclosure requirements, such as requiring

businesses in regulated professions to proactively disclose when consumers are interacting with AI.

Utah narrowed this framework through multiple bills in 2025, a reflection of early implementation concerns raised by state legislators and industry stakeholders. SB 226 and SB 332 narrowed the scope of disclosure obligations, limiting these obligations to "clear and unambiguous" consumer requests or "high-risk" interactions involving sensitive data and consequential advice, and narrowing the definition of covered AI systems to exclude routine uses (such as technologies that do not simulate human communication or generate human-like, nonscripted outputs). The Utah Division of Consumer Protection has authority to enforce SB 149, though enforcement remains limited to date.

New York: RAISE Act

California's Transparency in Frontier AI Act (TFAIA) was one of the first state regulatory frameworks for developers of frontier models. As we discussed in this April 1 alert, New York has since revised its frontier AI framework to align more closely with California's law. New York Gov. Kathy Hochul signed the Responsible AI Safety and Education (RAISE) Act in December 2025 with the expectation that legislators would amend the law to mirror TFAIA. Hochul signed those amendments on March 27, 2026, shifting the RAISE Act toward a transparency and reporting-based framework.

As revised, the RAISE Act imposes:

- Model-level obligations, including transparency and reporting on training, deployment, safety protocols and incidents.
- A shift away from deployment restrictions, removing earlier prohibitions on models posing an "unreasonable risk of critical harm."
- Alignment with California's framework, emphasizing safety testing, documentation and reporting.
- Key differences from TFAIA include:
    - Higher civil penalties (up to $1 million for a first violation and up to $3 million for subsequent violations).
    - Shorter incident reporting timeline (72 hours versus TFAIA's 15-day window). Other states, including Utah and Illinois, are considering similar frontier model regulation.

Key takeaways for companies

The evolution of AI laws is not limited to the US. The European Union AI Act – one of the earliest and most comprehensive cross-sector AI laws, imposing obligations on AI models based on risk tiers and categories of models – is also now being reconsidered by EU lawmakers for revision. While the AI Act entered into force on August 1, 2024, key obligations were set to phase in over time, with the main requirements starting in 2026, and certain obligations extending into 2027. However, the European Commission's November 2025 "Digital Omnibus" proposal, now advancing through the legislative process, would delay application of certain high-risk AI requirements and make targeted changes to exemptions, governance and implementation. As of April 2026, EU institutions are actively considering pushing key compliance deadlines to 2027 – 2028, reflecting implementation challenges and concerns about regulatory burden. The EU's AI regulatory framework continues to be refined and tailored in real time.

In combination, these developments underscore a broader shift: Even the most comprehensive AI regulatory regimes are being recalibrated as implementation approaches. Given the pace of change to these regulations, companies may benefit from a phased approach to compliance that accounts for evolving requirements and still emerging enforcement priorities. As such, companies should consider the following:

1. **Reassessing compliance strategies:** Several key state AI laws have upcoming deadlines, but some requirements are subject to amendment or delay. While enforcement currently has been minimal, regulators may begin issuing guidance and early enforcement actions as compliance dates approach.
2. **Monitoring federal action:** With the recent release of its AI legislative recommendations, the White House outlined an innovation-oriented federal approach to AI, recommending Congress preempt state laws that relate to AI development, "unduly burden" lawful activity assisted by AI or "penalize AI developers" for unlawful third-party conduct. Even if Congress does not act, or does so slowly, the administration is positioned to move through executive and enforcement channels. The Department of Justice's AI Litigation Task Force is expected to identify and potentially challenge state AI laws in court, and other federal

agencies, such as the Department of Commerce, may target certain states regulating AI by restricting federal funds. As such, companies should monitor both congressional developments and near-term federal activity, as the administration considers multiple pathways to shape the AI regulatory landscape.
3. **Tracking state-level changes:** As state AI regulation evolves, key areas to watch include the revisions to Colorado SB 205, further changes to California's AI laws, and the continued development of frontier regulations, including in New York and potentially Illinois, Washington and Utah.

Note

1. ADMT is defined as an automated decision-making technology that processes personal data to generate outputs (including predictions, scores and classifications) and is used to materially influence a consequential decision. The definition excludes (i) basic web infrastructure (such as web hosting and caching) that require human analysis and do not use machine learning; (ii) tools that solely summarize, organize or present information for human review; and (iii) general-purpose technology that provides information or recommendations.

This content is provided for general informational purposes only, and your access or use of the content does not create an attorney-client relationship between you or your organization and Cooley LLP, Cooley (UK) LLP, or any other affiliated practice or entity (collectively referred to as "Cooley"). By accessing this content, you agree that the information provided does not constitute legal or other professional advice. This content is not a substitute for obtaining legal advice from a qualified attorney licensed in your jurisdiction, and you should not act or refrain from acting based on this content. This content may be changed without notice. It is not guaranteed to be complete, correct or up to date, and it may not reflect the most current legal developments. Prior results do not guarantee a similar outcome. Do not send any confidential information to Cooley, as we do not have any duty to keep any information you provide to us confidential. When advising companies, our attorney-client relationship is with the company, not with any individual. This content may have been generated with the assistance of artificial intelligence (AI) in accordance with our AI Principles, may be considered Attorney Advertising and is subject to our legal notices.

Key Contacts

Tracy Rubin **Palo Alto**	**trubin@cooley.com** **+1 650 843 5568**
Sean Quinn New York	squinn@cooley.com +1 202 728 7075
Chris Chynoweth New York	cchynoweth@cooley.com +1 650 843 5372
Isamar Vaquero Washington, DC	ivaquero@cooley.com +1 202 776 2217
Adam Silow New York	asilow@cooley.com +1 212 479 6163
Richard Koch Washington, DC	rkoch@cooley.com +1 202 776 2323

This information is a general description of the law; it is not intended to provide specific legal advice nor is it intended to create an attorney-client relationship with Cooley LLP. Before taking any action on this information you should seek professional counsel.

Copyright © 2023 Cooley LLP, 3175 Hanover Street, Palo Alto, CA 94304; Cooley (UK) LLP, 22 Bishopsgate, London, UK EC2N 4BQ. Permission is granted to make and redistribute, without charge, copies of this entire document provided that such copies are complete and unaltered and identify Cooley LLP as the author. All other rights reserved.