

CLIENT ALERT | October 5, 2026

Newest Slate of AI Laws Keeps California at the Forefront of US Regulation

During the 2026 legislative session, California enacted more than two dozen AI laws regulating transparency, chatbot safety, automated decision-making, and other domains.

Key Points:

- **Expanded AI transparency requirements.** Amendments to CAITA eliminate its 1 million user threshold, significantly expanding the scope of covered providers subject to its requirements (SB 1000), and clarify how large online platforms must enable provenance-data inspection (AB 2713).
- **Chatbot safety and consumer protection.** New laws implement sweeping child safety reform for operators of companion chatbots (SB 1119), a moratorium on chatbot-equipped toys (SB 867), and transparency requirements for customer service chatbots (AB 1609).
- **Employer ADMT.** California will require human corroboration, employee data access, and post-use notice for disciplinary or termination decisions made by automated decision-making (SB 947).
- **Executive action on frontier AI safety.** Executive Order N-9-26 directs state agencies to accelerate implementation of auditor-requirement deadlines and develop recommendations for stronger frontier-AI safeguards, including onsite auditors, kill-switch requirements, and expanded incident reporting.
- **Broad sectoral reach.** California enacted some 20 other AI-focused laws addressing employment, intellectual property, data privacy, healthcare, elections, antitrust, and more.

Since 2024, California has enacted dozens of laws governing artificial intelligence across a range of sectors and industries, cementing itself as the de facto leader of AI regulation in the US. It was little surprise then that this year brought more of the same, with lawmakers sending more than 30 AI-focused bills to Governor Gavin Newsom's desk. But the approach has shifted slightly: Whereas previous years saw California introduce many first-of-their-kind AI frameworks, this year's slate largely builds on existing regimes and fills gaps that have emerged as the technology rapidly evolves.

This Client Alert focuses on the most impactful developments across four areas: expanded AI-content transparency rules, chatbot safety and consumer-protection mandates, workplace restrictions on

automated decision systems, and executive action on frontier AI safety. It also briefly touches on a number of AI laws focused on other issues that will take effect after receiving the governor's signature.

California AI Transparency Act – Amendments Under SB 1000 and AB 2713

The California AI Transparency Act (CAITA), which just came into effect on August 2, 2026, establishes baseline provenance and disclosure requirements for generative AI systems in California. As enacted, CAITA required covered providers with more than 1 million monthly visitors or users to make available a free “AI detection tool” that enables users to determine whether content was generated by the provider’s model; required covered providers to include latent disclosures in AI-generated image, video, and audio content; and, beginning January 1, 2027, would prohibit large online platforms from knowingly stripping provenance data from content. Just months after the law went into effect, California has now made several major amendments to CAITA’s scope and mechanics via Senate Bill 1000 (SB 1000) and Assembly Bill 2713 (AB 2713).

SB 1000 — Amendments to Covered Provider Scope and Disclosure Requirements

Removal of the User Threshold

SB 1000’s most consequential change is the deletion of the 1 million monthly visitor/user threshold from the definition of “covered provider.” Under the original CAITA, only persons that create, code, or otherwise produce a generative AI (GenAI) system with over 1 million monthly visitors or users and that is publicly accessible in California qualified as covered providers. By eliminating this threshold, SB 1000 makes any person producing a publicly accessible GenAI system within the state subject to CAITA’s requirements. This will substantially increase the number of businesses subject to CAITA’s obligations, expanding to include startups, open-source model providers, smaller SaaS platforms, and enterprise-facing tools that may not have reached the user minimum as the law was originally written.

Disclosure Verification Tool

CAITA already requires covered providers to offer a free tool that allows users to assess whether image, video, or audio content was created or altered by the provider’s GenAI system. SB 1000 renames the tool as a “disclosure verification tool” and makes several substantive changes to its requirements.

1. SB 1000 replaces the prior blanket prohibition on outputting personal provenance data with a consent-based framework. The tool may now output personal information if the affected user expressly consents, after receiving notice about what will be output and that the information becomes a permanent part of the file’s digital footprint.
2. SB 1000 permits covered providers to satisfy the tool requirement by directing users to a compliant third-party tool, provided it is compatible with the provider’s latent disclosures and clearly accessible through the provider’s interface.

3. The tool's assessment scope now excludes "minor modifications," such as cropping, file resizing, color changes, and format conversions.

SB 1000 also expands a covered provider's ability to restrict access to the tool. Under CAITA, covered providers could impose "reasonable limitations on access" to their tool in order to prevent demonstrable risks to the security or integrity of their AI systems. SB 1000 maintains that restriction but also allows covered providers to further restrict access to the tool in order to prevent its misuse for "malicious purposes."

Latent Disclosures

SB 1000 deletes the requirement that covered providers offer users the option to include a manifest disclosure in AI-generated content. For content disclosures, covered providers now need to support only latent disclosures. However, SB 1000 expands the scope of CAITA's latent disclosure requirement from content created by the GenAI system to content created *or altered* by the system, unless the alteration constitutes a "minor modification."

CAITA already required such latent disclosures to convey (A) the provider's name, (B) the name and version number of the GenAI system, (C) the time and date of creation or alteration, and (D) a unique identifier. SB 1000 alters the obligation to disclose the model's "version number" to instead require "version information," which is at least ostensibly a broader category, although SB 1000 provides no detail on what, if any, new information this disclosure should include. It also adds two new elements: (E) whether the GenAI system created or altered the content, which helps differentiate fully AI-generated content from AI-altered content; and (F), beginning January 1, 2029, whether the GenAI system is designed to primarily function as assistive technology.

Assistive Technology Carve-Out

SB 1000 delays the application of CAITA to GenAI systems designed to primarily function as "assistive technology" until January 1, 2029. SB 1000 also adds a new definition of "assistive technology," covering an item, piece of equipment, or product system used to increase, maintain, or improve the functional capabilities of individuals with disabilities, as well as services that directly assist an individual with a disability in selecting, acquiring, or using such an item, equipment, or product system. Covered providers are prohibited from falsely representing that a GenAI system is designed to primarily function as assistive technology, with violations subject to a civil penalty of \$50,000 per violation, and with each day of noncompliance treated as a discrete violation.

Enforcement

CAITA's existing enforcement framework carries forward under SB 1000, including a \$5,000 civil penalty per violation that is enforceable by the Attorney General, a city attorney, or a county counsel. Each day that a violation continues constitutes a discrete violation.

Notably, a covered provider that falsely represents that its system is designed to primarily function as assistive technology is subject to a standalone \$50,000-per-violation penalty until January 1, 2029.

SB 1000 adds a retroactivity safe harbor so that civil actions filed before SB 1000's effective date cannot be maintained if the alleged conduct does not violate CAITA on and after that date.

Effective Date

SB 1000 is an urgency statute, meaning it took effect immediately upon being signed into law on September 30, 2026.

AB 2713 – System Provenance Data Inspection

CAITA requires large online platforms — defined as public-facing social media or other similar platforms with more than 2 million unique monthly users that distribute content to users who did not create such content — to detect provenance data in content, provide a user interface disclosing the availability of system provenance data, and to allow users to inspect available provenance data in an easily accessible manner.

AB 2713 clarifies that a large online platform may satisfy the inspection requirement by any of three means: (1) displaying the system provenance data directly through the platform's user interface; (2) providing a link to a website or other application that displays the system provenance data, including a website or application operated by a third party; or (3) allowing a user to download any provenance data embedded into, attached to, or otherwise associated with the content (subject to applicable federal copyright laws) in a format that cannot easily be embedded into, attached to, or associated with unrelated content. The law also specifies that the large online platform must not knowingly strip system provenance data or digital signatures from content uploaded to, distributed on, or downloaded from the platform, to the extent technically feasible. AB 2713 goes into effect on January 1, 2027.

Strategic Implications for Developers

For companies that already fell under CAITA's scope before these amendments, the path forward is largely incremental. SB 1000 will require those companies to add the two new latent-disclosure data fields (with the assistive-technology indicator deferred to 2029), rebrand the detection tool as a "disclosure verification tool," and, if desired, remove any manifest-disclosure functionality that is no longer required.

The real compliance challenge falls on the wave of newly covered providers, which will likely now include startups, open-source model developers, enterprise-facing SaaS platforms, and other GenAI system operators that previously fell below the 1 million-user threshold. For these companies, the immediate priority should be building or licensing compliant disclosure verification tools, implementing latent-disclosure pipelines in their content-generation infrastructure, and setting up reliable notification and noncompliance-reporting processes. Companies currently building GenAI tools should immediately incorporate CAITA compliance into product design rather than attempting to retrofit at scale.

Large online platforms should separately evaluate whether their current capabilities satisfy both AB 2713's revised user-interface requirements and at least one of its three inspection-compliance options (direct UI display, a link to a first- or third-party display site, or downloadable provenance data in a non-

transferable format). Platforms should also implement technical safeguards to ensure provenance data and digital signatures are not stripped during upload, distribution, or download.

More broadly, the elimination of CAITA's user threshold signals California's intent to continue expanding its regulatory footprint even within areas where it had already seemingly taken a leading position. Companies should be prepared for lawmakers to continue building on other existing regimes.

Chatbot Regulation – SB 1119, SB 867, and AB 1609

California had previously enacted a number of chatbot laws that have served as benchmarks for other states to follow. The 2026 session continued this trend, with lawmakers adding several new significant user-safety and consumer-protection obligations for chatbot operators in the state.

SB 1119 – Companion Chatbots: Children's Safety (Adam's Law)

SB 1119, known as "Adam's Law," establishes a comprehensive child safety framework focused on protecting minors from potential harms that could result from interaction with companion chatbots, with most substantive requirements going into effect on July 1, 2027.

Companion Chatbot Definition

The law defines a "companion chatbot" as an AI system with a natural language interface that provides adaptive, human-like responses to user inputs and is capable of meeting a user's social needs, including by exhibiting anthropomorphic features and sustaining a relationship across multiple interactions.¹

The law applies to "operators" of companion chatbots, defined as any person that makes a companion chatbot available to users in California. Notably, the "operator" definition excludes companies that make chatbots available exclusively to employees, contractors, or other personnel for use in a workplace setting, or postsecondary educational institutions that use these chatbots for educational purposes.

Age Assurance

SB 1119 requires operators to either verify user age pursuant to the Digital Age Assurance Act (DAAA) or the California Health and Safety Code, or to otherwise apply child protections to all users by default.

Under the DAAA, operating system providers must create an interface that requires an account holder to indicate the user's birth date or age during account setup. The system provider then generates "age bracket data" via a digital signal transmitted through a real-time application programming interface (API) indicating the user's age bracket. Developers must then request the signal from the operating system provider with respect to users that download and launch their application, at which point the developer is deemed to have actual knowledge of the user's age range.

If an operator cannot verify age through the DAAA signal, it can instead rely on an age determination made pursuant to California Health and Safety Code section 27001(a)(1)(B). Under this provision, operators may use one or more commercially reasonable age-assurance methods that are reasonably

effective, measurably consistent, and testable, such as biological or behavioral signals, verified user information, cryptographic techniques, government-issued identification (though not as the sole method), or an ISO/IEC 27566-compliant method. Operators may not rely on self-declaration, general contractual age restrictions, or payment methods available to minors.

As an alternative compliance path, SB 1119 also permits operators to apply the “Child Safety Measures” and “Data and Advertising Restrictions” described below to all users by default.

Risk Assessment

Before making a new or substantially modified companion chatbot available to users in California, operators must perform a comprehensive risk assessment with respect to children (defined as any person under the age of 18). This assessment must include a summary of the operator’s evaluation of potential “covered harms” associated with the use of the chatbot, which include reasonably foreseeable physical or financial harm; severe and reasonably foreseeable psychological or emotional harm to a reasonable child; a highly offensive intrusion on privacy rights protected by state or federal law; or adverse discrimination in violation of state or federal law. The assessment must also include a high-level description of the methodology used to evaluate covered harms. Operators must take and document measures that would reasonably mitigate any child safety risks identified in the risk assessment.

If another law requires the operator to perform a risk assessment that is substantially similar in scope, the operator may use that risk assessment to comply with SB 1119. An operator that controls a “comparable set” of companion chatbots only needs to perform a single risk assessment for the set.

Child Safety Policy

Operators that permit children to use their companion chatbots must publish on their website a child safety policy that describes how the chatbot is designed to prevent covered harms, as well as how the operator complies with SB 1119’s various requirements, including age assurance, data and advertising restrictions, default settings, and crisis response protocols.

Operators that prohibit children from using their companion chatbots must publish on their website a description of how the operator complies with its obligation to conduct age-assurance checks under SB 1119.

Child Safety Measures

SB 1119 will require operators that allow children to access their companion chatbots to implement a series of child safety measures, including:

- A documented crisis response protocol (including, among other things, timely referrals to crisis resources and parental notification if the child’s account is linked to their parent’s account), as well as a duty to preserve conversations between a child user and chatbot that triggered such protocol for at least three years

- Usage reminders, age-appropriate risk disclosures, and protective design features related to any identified child safety risks
- Certain default settings changeable only by parents (including disabling persistent conversational memory and push notifications, single-session limits of one hour, and daily limits of two hours)
- Periodic AI-disclosure notices during extended interactions (though no specific interval for such notices is established by the law)
- Reasonable measures to prevent the chatbot from engaging in potentially harmful behaviors, such as encouraging self-harm, diagnosing health conditions, engaging in or producing sexual content, claiming sentience, soliciting gifts or in-app purchases, expressing romantic interest, encouraging reliance for emotional support, using excessive flattery, or helping circumvent parental controls
- An interface design that allows children and parents to understand and utilize any protective features and controls
- A public incident-reporting mechanism that allows anyone to report incidents regarding child safety risks to the operator

Data and Advertising Restrictions

Operators may not display cross-context behavioral advertising to children, target advertising at a child using personal information in chat, sell child users' personal information, or use dark patterns with child-facing controls.

Child Safety Audits

By January 1, 2029, operators must complete an independent child safety audit, with subsequent audits occurring every two years thereafter (or before releasing a substantial modification that increases child safety risk). After receiving a report from an independent auditor, the operator must submit an audit summary to the California Attorney General within 30 business days and post a high-level summary on its website within 90 calendar days. Operators with less than \$500 million in gross annual revenue are exempt from audit requirements until January 1, 2032.

Enforcement

Public prosecutors may bring civil actions seeking penalties up to \$5,000 per affected child for negligent violations, or \$15,000 per affected child for intentional violations. A child (or a parent or guardian acting on that child's behalf) who suffers actual harm from a violation of the law's child safety measures may also bring a private action for actual damages, injunctive or declaratory relief, and reasonable attorney's fees.

SB 867 – Toys with Companion Chatbots

Senate Bill 867 (SB 867) amends California's existing companion chatbot law to prohibit the manufacture, sale, or possession with intent to sell or exchange any toy that includes a companion chatbot. The term

“companion chatbot” uses the same definition as SB 1119, while a “toy” is defined as any physical product designed, marketed, or manufactured for use in play by children under 16 years of age. The prohibition takes effect on January 1, 2027, and is repealed on January 1, 2031, creating a four-year moratorium.

AB 1609 – Customer Service Chatbots

Assembly Bill 1609 (AB 1609) regulates customer service chatbots used by “large private businesses,” defined as businesses with gross annual revenue over \$500 million that provide goods and services to California consumers.

Disclosure Requirements

Large private businesses using customer service chatbots may not affirmatively represent that the chatbot is human. If a reasonable person would be misled into believing the chatbot is human, the business must provide a clear and conspicuous disclosure to users indicating that the chatbot is AI.

Access to Human Agent

During regular business hours, a large private business must provide customers with a feature that allows them to request to be connected to a human customer service agent. The business must make a good faith effort to connect the customer to a human agent within 15 minutes or schedule a specific appointment time within one business day of the request. A large private business that provides goods and services to California customers through online platforms and that also maintains a telephonic customer service platform must post the customer service number clearly and conspicuously on its website.

Enforcement

AB 1609 is enforceable by public prosecutors, with penalties of up to \$5,000 for an initial violation and \$10,000 for subsequent violations.

Exemptions

The bill waives requirements during unforeseen circumstances or extraordinary/emergency situations (such as statewide emergencies or power shutoffs), and exempts exclusive business lines and communications by hospitals and consumer reporting agencies.

Strategic Implications for Chatbot Operators

The convergence of SB 1119 and SB 867 with California’s existing companion chatbot laws (such as SB 243) creates a layered regulatory architecture that will require companion chatbot operators to rethink product design, safety infrastructure, and market strategy. In particular, SB 1119 will require that child safety be engineered into the product well before launch, with documented risk assessments; documented crisis response protocols that provide referrals to crisis services and, where a credible and imminent threat is identified, either parental notification or streamlined access to the 988 Suicide and Crisis Lifeline or an equivalent crisis helpline; and default usage limits that can only be adjusted by parents.

Operators should begin scoping these technical requirements as soon as practicable, particularly the age-assurance integration with the DAAA, which will require coordination with operating system providers and application stores to implement the real-time API-based age signals that the statute contemplates. The biennial independent audit requirement, although not applicable until January 1, 2029, should not be pushed to the back burner either. There will be significant lead time required to identify qualified third-party auditors, negotiate engagement terms, and build the internal documentation and control frameworks that auditors will need in order to provide a compliant audit report.

For large consumer-facing businesses, AB 1609's 15-minute human-connection mandate may be the most challenging new requirement from an operations perspective, particularly for companies that have already invested in automated customer service systems to reduce costs. These businesses should start to model staffing, infrastructure, and scheduling needs that will result from guaranteeing human agent access during regular business hours. While AB 1609's "good faith effort" standard will certainly provide some flexibility to companies that cannot always meet the 15-minute timeframe, the limits of that standard will almost certainly be tested through enforcement that will help set expectations once the law goes into effect.

Automated Decision-Making Technology – SB 947

Scope

SB 947 regulates employer use of automated decision systems in the employment context. It defines an "automated decision system" (ADS) as a computational process derived from machine learning, statistical modeling, data analytics, or artificial intelligence that issues simplified output, including a score, classification, or recommendation, that is used to assist or replace human discretionary decision-making and materially impacts natural persons. The definition expressly excludes spam filters, firewalls, antivirus software, identity and access management tools, calculators, databases, datasets, and other compilations of data. SB 947 also introduces the concept of "ADS output," which is any information, data, assumptions, predictions, scoring, recommendations, decisions, or conclusions generated by an ADS.

Prohibitions

An employer cannot use an ADS to: (1) prevent compliance with or violate any federal, state, or local labor, occupational health and safety, employment, or civil rights laws or regulations; (2) infer an employee's protected status under the California Fair Employment and Housing Act; or (3) predict and take adverse action against a worker for exercising legal rights.

Human Corroboration Requirement

SB 947 prohibits employers from relying solely on ADS outputs to make disciplinary or termination decisions. If an employer intends to primarily rely on ADS for such a decision, the employer must have a human corroborate the decision either by manually reviewing the data used by the ADS to generate the decision, or by reviewing other relevant supporting information such as supervisory evaluations, personnel files, work product, peer reviews, or witness interviews. If the human reviewer cannot

corroborate the ADS output or finds it inaccurate, incomplete, or misleading, the employer cannot use the ADS output as part of the decision-making process.

Employee Data Access

Upon request by an affected employee, an employer must provide a meaningful, objective description of the employee's own data used by the ADS to make the disciplinary or termination decision. Any description provided by the employer must be anonymized to protect personal information belonging to customers, employees, or any other individuals.

Post-Use Notice

When an employer primarily relies on ADS to make a disciplinary or termination decision, the employer must provide the affected employee with separate written notice at the time it informs the employee of the decision. This separate notice must (1) tell the employee that the employer primarily relied on ADS to make the decision, (2) confirm that a human reviewed and corroborated the ADS decision, (3) provide contact information for a human that the employee can contact for more information about the decision and the employee's right to access data used by the ADS, and (4) note that the employer is prohibited from retaliating against the employee for exercising their rights.

Enforcement and Exemptions

Violations are enforceable by the Labor Commissioner or through a civil action brought by a public prosecutor. Employers may not fire, demote, suspend, or otherwise retaliate against any affected employee for filing a complaint with the Labor Commissioner or cooperating in an investigation of an alleged violation. Employers are subject to a civil penalty of \$500 per violation, as well as injunctive relief, punitive damages, and attorney's fees.

Employers that comply with SB 947's notice requirements are not required to comply with substantially similar notice provisions under other California law relating to the use of ADS in employment. However, employers subject to the California Consumer Privacy Act (CCPA) remain subject to privacy-related automated decision-making technology regulations adopted by the California Privacy Protection Agency (CPPA).

The law exempts parties covered by a collective bargaining agreement that explicitly waives the statute in clear and unambiguous terms and that provides protection from algorithmic management. It also exempts the use of ADS where required by a federal statute, regulation, or contract for national security or for the development of aircraft used in national airspace.

Effective Date

SB 947 will go into effect on July 1, 2027.

Strategic Implications for Employers

SB 947's relatively narrow focus on disciplinary and termination decisions means it will not regulate every employer use of AI, but the decisions it does cover tend to be among the most litigation-prone in employment law. Employers that use AI-driven performance monitoring, productivity scoring, or conduct-flagging tools must develop a clear understanding of which of those systems produce outputs that feed into disciplinary or termination decisions within the organization. Given the data rights and requirements established under the statute (including the requirement that human corroboration may rest on the data underlying the ADS output or on other relevant supporting information), employers must ensure any ADS systems are sufficiently transparent and explainable to permit them to isolate the underlying data for review.

SB 947's interaction with existing California privacy law is also worth a close analysis. While compliant employers are exempt from similar notice requirements under other state laws, employers subject to the CCPA remain independently obligated to comply with any ADMT regulations adopted by the California Privacy Protection Agency, meaning that any workflows should ideally be designed to satisfy both SB 947 and applicable CCPA/CPPA requirements in a single process.

Executive Order N-9-26

On September 18, 2026, Governor Newsom issued Executive Order N-9-26, directing the California Government Operations Agency (GOA) to take several steps that could further alter the state's AI regulatory regime.

Recommended Amendments to Frontier AI Laws

By November 16, 2026, GOA must submit recommendations on the technical feasibility of amending existing frontier AI laws to address (1) onsite IVO audits at frontier developer labs, (2) independent verification of safety frameworks and risk assessments, (3) "kill switch" requirements for frontier models, and (4) expanded critical safety incident reporting to include loss-of-control incidents.

Agency Implementation Deadlines

By May 1, 2027, GOA must complete the requirements of Government Code section 8898.1 by developing application requirements, procedures, and criteria for independent verification organizations (IVOs) and publicly posting them. Section 8898.1 was enacted as part of SB 813 (discussed in the Additional Legislative Developments section below), which creates a state framework for qualifying and overseeing IVOs that have demonstrated expertise in assessing the risks posed by AI systems or models. Notably, SB 813 calls for GOA to complete these tasks "on or before January 1, 2028," so the order accelerates GOA's outer deadline by eight months.

By December 1, 2027, the agency must also complete the requirements of Government Code section 11549.82(a) and begin taking action under subdivision (b), which was enacted as part of AB 1405 (discussed in the Additional Legislative Developments section below). That law requires GOA to create a mandatory AI auditor registry, set annual registration fees, establish a public mechanism for reporting

auditor misconduct, publish auditor information, and retain misconduct reports. As with SB 813, the order's December 2027 deadline would significantly speed up AB 1405's timeframe, as the law originally called for GOA to complete these actions by January 1, 2029.

Strategic Implications

The executive order does not directly impose compliance mandates on private companies, but GOA's November 2026 recommendations could lay the groundwork for future legislation that would introduce major new obligations for frontier developers. While GOA's recommendations would still need to go through the standard legislative process in order to become law, frontier developers should monitor these recommendations closely to stay ahead of any potential new developments.

The order's acceleration of implementation timelines for SB 813 (IVO framework, now due May 1, 2027) and AB 1405 (AI auditor registry, now due December 1, 2027) means that the state's third-party verification and audit infrastructure will come online faster than the underlying statutes originally contemplated. Companies that will be subject to audit requirements under California law should likewise follow the application criteria and procedural requirements that GOA establishes in the coming months.

Additional Legislative Developments

In addition to those discussed above, Governor Newsom signed a suite of new AI-focused bills into law that cover a wide range of topics and issues. The most notable new laws are briefly discussed below:

ADMT and AI Governance

- **AB 1405.** Requires GOA to establish an AI auditor registry that AI auditors must use to register before conducting AI audits necessary for compliance with California law. Also imposes standards of independence, objectivity, and integrity on registered auditors.
- **SB 813.** Requires GOA to establish criteria for designating IVOs to assess AI systems' risks and compliance, and to convene stakeholder working groups to develop standards. IVOs must submit annual reports. The law does not create an independent audit mandate for private companies.

Synthetic Media and Digital Replicas

- **SB 1050.** Makes it unlawful to create and publish an advertisement prominently including a synthetic performer without a clear and conspicuous disclosure. Violations constitute violations of California's false advertising statute.
- **SB 1111.** Amends California right-of-publicity and criminal false impersonation laws to cover the use of AI-generated digital replicas.
- **SB 1276.** Expands child sexual exploitation offenses to cover knowingly creating, exchanging, downloading, streaming, or accessing digitally altered or AI-generated matter depicting a person under 18 engaged in sexual conduct.

AI in Employment and Workforce

- **SB 951.** Amends California's Worker Adjustment and Retraining Notification (WARN) Act to require employers giving notice of mass layoffs caused by AI or automation to include specified technological displacement information. Requires the Employment Development Department to publish quarterly summaries of technology-related displacements, as well as a legislative report by January 1, 2028.
- **AB 1883.** Prohibits employers from using workplace surveillance tools that use AI to recognize or infer emotional states or collect neural data, with exceptions for national security and defense operations. Employers face penalties of up to \$500 per violation.
- **AB 1331.** Prohibits employers from using workplace surveillance tools to monitor or surveil employees in bathrooms and grants employees the right to leave behind workplace surveillance tools prior to entering a bathroom.

Data Privacy

- **SB 354.** Overhaul of California insurance privacy law establishing new standards for processing and sharing consumers' personal information by insurance licensees, reinsurers, and service providers. With respect to AI, the law defines "personal information" to include data existing in various formats, including AI systems capable of outputting personal information.

AI Transparency

- **SB 1159.** Clarifies that "person," "member of the public," and similar terms in California's major transparency and governance laws do not include AI systems, and prohibits knowingly using AI to falsely represent human engagement with governmental agencies.

AI in Healthcare

- **SB 503.** Requires developers and deployers of AI clinical decision-support systems to identify and mitigate risks of bias and further requires developers to provide documentation to deployers regarding intended uses and risks of such systems.
- **AB 1979.** Requires health facilities to ensure licensed providers retain independent professional judgment when care is informed by AI output, and prohibits providers from using AI to independently perform, or direct unlicensed staff in performing, clinical functions that require a license. Deems businesses offering healthcare chatbots to be providers subject to the Confidentiality of Medical Information Act.

AI in Education

- **SB 928.** Requires California State University (CSU) instructors of record to meet the trustees' faculty-qualification rules, effectively preventing AI from serving as the instructor, though not prohibiting employees from using AI tools in their work.

- **AB 2392.** Requires the California Community Colleges and CSU (and requests the University of California) to convene a working group that will develop GenAI procurement standards and training by January 1, 2028, and to report systemwide GenAI contracts to the Legislature in the meantime.
- **AB 1159.** Amends student privacy laws and enacts the Higher Education Student Information Protection Act to prohibit use of student information to train GenAI or develop AI systems, among other things.

AI in Elections

- **AB 502.** Prohibits the knowing, malicious distribution of materially deceptive digitally created or modified content (including deepfakes) that is material to voters' electoral decisions or public confidence in elections, within specified windows before and after elections. This is an urgency statute and took effect when signed into law.

AI in Legal Practice

- **SB 574.** Prohibits attorneys from delegating the practice of law and arbitrators from delegating decision-making to GenAI; requires verification of AI-generated outputs, disclosure of GenAI use to courts, and protection of confidential information when using GenAI.

Contacts

[Michael H. Rubin](#)
michael.rubin@lw.com
+1.415.395.8154
San Francisco

[Andrew Gass](#)
andrew.gass@lw.com
+1.415.395.8806
San Francisco

[Ghaith Mahmood](#)
ghaith.mahmood@lw.com
+1.213.891.8375
Los Angeles

[Sy Damle](#)
sy.damle@lw.com
+1.202.637.3332
+1.212.906.1659
Washington, D.C. / New York

[Fiona Maclean](#)
fiona.maclean@lw.com
+44.20.7710.1822
London

This publication is produced by Latham & Watkins as a news reporting service to clients and other friends. The information contained in this publication should not be construed as legal advice. Should further analysis or explanation of the subject matter be required, please contact the lawyer with whom you normally consult. The invitation to contact is not a solicitation for legal work under the laws of any jurisdiction in which Latham lawyers are not authorized to practice. See our [Attorney Advertising and Terms of Use](#).

Endnotes

-
- ¹ The definition is adopted from SB 243, another companion chatbot law enacted by California in 2025. The definition excludes customer-service bots, video-game bots limited to game topics, and standalone voice-assistant devices. An “operator” is a person who makes a companion chatbot available in the state.