
UNFINISHED BUSINESS

Bill No: SB 1000
Author: Becker (D)
Amended: 8/21/26
Vote: 27 - Urgency

SENATE PRIV., DIGITAL TECH. & CONS. PROT. COMMITTEE: 8-1, 4/13/26
AYES: Cabaldon, Gonzalez, McNerney, Ochoa Bogh, Padilla, Reyes, Umberg,
Wiener
NOES: Jones

SENATE APPROPRIATIONS COMMITTEE: 5-0, 5/14/26
AYES: Cervantes, Cabaldon, Grayson, Richardson, Wahab
NO VOTE RECORDED: Seyarto, Dahle

SENATE FLOOR: 33-1, 5/19/26
AYES: Allen, Archuleta, Arreguín, Ashby, Becker, Blakespear, Cabaldon,
Caballero, Cervantes, Cortese, Durazo, Gonzalez, Grayson, Grove, Hurtado,
Laird, Limón, McGuire, McNerney, Menjivar, Ochoa Bogh, Padilla, Pérez,
Reyes, Richardson, Rubio, Smallwood-Cuevas, Stern, Umberg, Valladares,
Wahab, Weber Pierson, Wiener
NOES: Choi
NO VOTE RECORDED: Alvarado-Gil, Dahle, Jones, Niello, Seyarto, Strickland

ASSEMBLY FLOOR: 72-0, 8/25/26 - See last page for vote

SUBJECT: California AI Transparency Act

SOURCE: Author

DIGEST: This bill makes a number of changes to the California AI Transparency (CAIT) Act, including eliminating the requirement to offer manifest disclosures, removing the user threshold on the definition of “covered provider,” amending the requirements applicable to third-party licensees, and amending other definitions and exemptions.

Assembly Amendments refine privacy protections, specify that various disclosure verification tool requirements may be fulfilled through the use of third-party disclosure verification tools, as provided, require latent disclosures to specify whether a generative artificial intelligence system is designed to primarily function as assistive technology, prohibit making false representations, among other changes.

ANALYSIS:

Existing law:

- 1) Establishes the CAIT Act, which becomes operative, in part, on August 2, 2026, and requires certain “covered providers” to make an artificial intelligence (AI) detection tool available at no cost by which a person can assess whether content was created or altered by the provider’s GenAI system. (Business (Bus.) & Professions (Prof.) Code § 22757 et seq.)
- 2) Requires a “covered provider,” a person that creates, codes, or otherwise produces a GenAI system that has over 1 million monthly visitors or users and is publicly accessible within the geographic boundaries of the state, to offer users the option to include in AI-generated image, video, or audio content created by its own generative AI system a manifest disclosure that meets specified criteria, including that it identifies the content as AI-generated content. (Bus. & Prof. Code § 22757.3(a).)
- 3) Requires a covered provider to include in AI-generated image, audio, and video content created by its generative AI system a latent disclosure that is detectable by the tool specified above and is, to the extent technically feasible, permanent or extraordinarily difficult to remove. (Bus. & Prof. Code § 22757.3(b).)
- 4) Prohibits a covered provider from doing any of the following in carrying out the duties above:
 - a) Collect or retain personal information when a person utilizes the covered provider’s AI detection tool, except that it may collect and retain the contact information of a person who submitted feedback.
 - b) Retain any content submitted to the AI detection tool for longer than is necessary to comply with this law. (Bus. & Prof. Code § 22757.2(c).)
- 5) Requires a large online platform, starting January 1, 2027, to do one of the following:

- a) Detect whether any provenance data that is compliant with widely adopted specifications adopted by an established standards-setting body is embedded into or attached to content distributed on the large online platform.
 - b) Provide a user interface to disclose the availability of system provenance data that reliably indicates that the content was generated or substantially altered by a GenAI system or captured by a capture device. The user interface shall make clearly and conspicuously available to users information sufficient to identify the content's authenticity, origin, or history of modification, including specified information such as whether provenance data is available.
 - c) Allow a user to inspect all available system provenance data that is compliant with widely adopted specifications adopted by an established standards-setting body in an easily accessible manner by any of several specified means. (Bus. & Prof. Code § 22757.3.1.)
- 6) Provides that violators of the above provisions are liable for a civil penalty in the amount of \$5,000 per violation to be collected in a civil action filed by the Attorney General, a city attorney, or a county counsel. Each day in violation is deemed a discrete violation. (Bus. & Prof. Code § 22757.4.)
- 7) Provides that the CAIT Act does not apply to any product, service, website, or application that provides exclusively non-user-generated video game, television, streaming, movie, or interactive experiences.

This bill:

- 1) Removes the user threshold from the definition of “covered provider.”
- 2) Removes the requirement for covered providers to offer users the option to include a manifest disclosure in content. Amends the latent disclosure requirement only to the extent “technically feasible.”
- 3) Redefines “provenance data” to mean information about the origin of a piece of digital content and the history of modifications to the content that is in a format that is compliant, or interoperable with, widely adopted specifications adopted by an established standards-setting body.
- 4) Replaces references to “AI detection tool” with “disclosure verification tool.”
- 5) Provides that the tool need not detect altered content if only a “minor modification,” which includes any of the following alterations:

- a) A change to brightness, contrast, or color.
 - b) Sharpening.
 - c) Saturating.
 - d) File resizing.
 - e) Scaling.
 - f) Cropping.
 - g) File format conversions.
 - h) Denoising and removal of background noise in audio.
- 6) Authorizes covered providers to impose reasonable limitations on access to the tool to prevent misuse of the tool for malicious purposes.
- 7) Reworks the provisions regarding a users' personal information. A covered provider shall not make access to the covered provider's GenAI system or disclosure verification tool contingent upon providing personal information beyond what is strictly necessary for the specified purposes.
- 8) Reworks third-party licensing provisions. Permits covered providers to notify a noncompliant licensee instead of immediately terminating their license. If licensee fails to cure violation and report back to covered provider, requires covered provider to notify Attorney General.
- 9) Makes other technical and clarifying changes to the CAIT Act.
- 10) Includes an urgency clause.

Background

Certain forms of media – audio recordings, video recordings, and still images – can be powerful evidence of the truth. While such media have always been susceptible to some degree of manipulation, fakes were relatively easy to detect. The rapid advancement of AI technology, specifically the wide-scale introduction of GenAI models, has made it drastically cheaper and easier to produce synthetic content created, audio, images, text, and video recordings that are not real, but that are so realistic that they are virtually impossible to distinguish from authentic content.

Last session, the Legislature passed the California AI Transparency Act (CAIT), which, beginning August 2, 2026, regulates provenance data disclosure in AI-generated content. “Covered providers” are required to embed latent provenance disclosures in content generated by their GenAI systems and to offer manifest disclosures. Providers are required to make available an AI detection tool at no cost to the user that meets certain criteria, including that it allows a user to assess whether certain content was created or altered by the covered provider’s system.

This bill is a product of many stakeholder conversations that seeks to clean up the CAIT Act and ensure it accomplishes its intended goal of providing useful transparency to Californians. The bill is author-sponsored. It is supported by Transparency Coalition.AI and ForensicVB LLC. No timely opposition was received.

Comment

According to the author:

As AI technology advances, distinguishing between human and machine-generated content becomes increasingly challenging. This ambiguity poses significant risks to Californians, exacerbating problems of disinformation, harassment, and fraud while threatening the integrity of the information environment our democracy and economy depend on. In 2024, the legislature passed SB 942, the first bill in the nation to establish disclosure requirements for synthetic content. Since then, content provenance technology has developed in such a way that the leading technologies for embedding provenance are not accurately described by the law. Additionally, some methods of marking content are still in their nascency, and aren’t yet robust enough to withstand adversarial exploitation or certain kinds of transformations (e.g., file compression, spoofing attacks, added noise to an image, or file type conversions). The original act also did not address the risk that visible labels create a binary signal implying that unlabeled content is authentic and labeled content is suspect - regardless of whether either inference is warranted. Furthermore, obligations for developers established by SB 942 do not require that the information embedded in AI generated or modified content be consistent - or interoperable with - widely adopted standards. Last year, AB 853 (Wicks) established obligations for large online platforms to read provenance data, but only if it was compliant with such standards. This gap risks establishing a fragmented ecosystem

where new methods of provenance disclosure that aren't interoperable with current standards are not read by large online platforms, and never subsequently relayed to users. In light of the low implementation cost of current provenance standards relative to costs of model development, this bill removes the user threshold from the definition of "covered provider" to expand obligations to any AI system which is publicly available within the state, and aligns the act's requirements to reflect the state of content provenance standards and international provenance disclosure regimes. This bill is critical to ensuring California's content labeling laws are effective at providing consumers with consistent information about where the content they see online comes from.

FISCAL EFFECT: Appropriation: No Fiscal Com.: Yes Local: No

According to the Senate Appropriations Committee:

- Unknown costs to the state funded trial court system (Trial Court Trust Fund, General Fund) to adjudicate civil actions. By expanding who may be considered a covered provider and, consequently, subject to civil penalties with statutory damages, this bill may lead to additional case filings. The fiscal impact of this bill to the courts will depend on many unknowns, including the number of cases filed and the factors unique to each case. An eight-hour court day costs approximately \$10,500 in staff workload. While the courts are not funded on a workload basis, an increase in workload could result in delayed court services and would put pressure on the General Fund to fund additional staff and resources and to increase the amount appropriated to backfill for trial court operations.
- The Department of Justice (DOJ) does not anticipate a significant fiscal impact.
- The California Department of Technology (CDT) notes that it would not fall within the definition of covered provider under this bill; therefore, it anticipates only minor and absorbable costs to continue tracking developments in the AI space.

According to the Assembly Appropriations Committee:

- Potential minor costs to the Department of Justice (General Fund) for enforcement. The Attorney General's authority to enforce the chapter already exists under current law, which this bill does not amend: a violation is

subject to a civil penalty of \$5,000 per violation, collected in a civil action filed by the Attorney General, a city attorney, or a county counsel. The bill does not create any new enforcement authority or mechanism. By removing the requirement that a covered provider have more than 1,000,000 monthly users, the bill enlarges the population of providers subject to the chapter, but any additional enforcement activity, and the associated costs, would likely be minor and would depend on whether the Attorney General elects to pursue actions under existing authority.

- The bill creates no new state program or appropriation, and the operative obligations apply to private covered providers. Enforcement by a city attorney or county counsel is permissive and would fall to local funds, so the bill imposes no reimbursable state-mandated local program. Any penalty revenue is not a significant or reliable funding source.

SUPPORT: (Verified 8/25/26)

Center for Ai and Digital Policy (CAIDP)

Forensicvb LLC

Transparency Coalition.ai

OPPOSITION: (Verified 8/25/26)

None received

ARGUMENTS IN SUPPORT: ForensicVB LLC writes:

From an operational standpoint, manifest disclosure or visible labels introduce a significant risk of investigative de-prioritization. When content is labeled—accurately or not—as AI-generated, it can create uncertainty at the front end of an investigation, delay triage decisions, and divert limited law enforcement resources away from time-sensitive CSAM cases. Put more simply, it has been reported that pedophiles operating on the Dark Web are labeling real CSAM content as “AI-generated” in the expectation that investigators will see the label and focus their attention on other non-labeled content.

For these reasons, I believe SB 1000’s current reliance on latent, provenance-based disclosure frameworks reflects a more responsible and technically sound path forward, and is consistent with existing content-provenance approaches already in use, such as the Coalition for Content Provenance and Authenticity (C2PA) specifications. The same professional experience that has led me to oppose manifest disclosure requirements has also led me to support latent

disclosure approaches as a more reliable, durable, and effective trust signal for consumers of digital content. When implemented correctly, latent disclosure mechanisms can support authenticity, attribution, and accountability without creating avoidable investigative friction or opportunities for misuse.

ASSEMBLY FLOOR: 72-0 8/25/26

AYES: Addis, Aguiar-Curry, Ahrens, Alanis, Alvarez, Arambula, Ávila Farías, Bains, Bauer-Kahan, Bennett, Berman, Boerner, Bonta, Bryan, Calderon, Caloza, Carrillo, Chen, Connolly, Davies, DeMaio, Dixon, Ellis, Flora, Fong, Gabriel, Garcia, Gipson, Jeff Gonzalez, Mark González, Haney, Harabedian, Hart, Hoover, Irwin, Jackson, Johnson, Kalra, Krell, Lee, Lowenthal, Macedo, Muratsuchi, Nguyen, Ortega, Pacheco, Papan, Patel, Patterson, Pellerin, Petrie-Norris, Quirk-Silva, Ramos, Ransom, Celeste Rodriguez, Michelle Rodriguez, Rogers, Blanca Rubio, Sanchez, Schiavo, Schultz, Sharp-Collins, Solache, Soria, Stefani, Valencia, Wallis, Ward, Wicks, Wilson, Zbur, Rivas

NO VOTE RECORDED: Castillo, Elhawary, Hadwick, Lackey, McKinnor, Ta, Tangipa

Prepared by: Christian Kurpiewski / P., D.T., & C.P. / (916) 651-1548, Benjamin Diaz / P., D.T., & C.P. / (916) 651-1548

9/14/26 10:40:49

**** END ****