THIRD READING

Bill No: SB 1000
Author: Becker (D)
Amended: 3/26/26
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

**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," and amending other definitions and exemptions.

**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. Requires the latent disclosure to be compliant or interoperable with widely accepted industry standards.

3) Redefines "provenance data" to mean information about the origin of a piece of 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) Resizing.
    e) Scaling.
    f) Cropping.
    g) Format conversions.
    h) Resampling.
    i) Denoising and removal of background noise in audio.

6) Permits the tool to output personal information that is detected in the content if the user to whom the personal information pertains expressly consents, clearly and conspicuously in plain language, to including personal information specified by the user in the content pursuant to a notice that does both of the following:

a) Informs the user of the personal information that may be output by the tool.
b) Informs the user that once personal information is embedded into provenance data and exported, the information becomes part of the file's permanent digital footprint and cannot be retracted from copies already in circulation.

7) Authorizes covered providers to impose reasonable limitations on access to the tool to prevent misuse of the tool for malicious purposes.

8) Reworks the provisions regarding a users' personal information to provide that a covered provider shall not collect, use, or retain personal information from a user of the covered provider's AI disclosure verification tool or any content submitted to the disclosure verification tool beyond what is reasonably necessary for user authentication. 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.

9) Reworks the third-party obligations to provide that a covered provider must require by contract that the third-party licensee ensures that the system includes provenance data that meets the specified criteria laid out in the CAIT Act to the extent it is technically feasible. It reduces the time within which a covered provider must revoke a license from 96 hours to 48 hours.

10) Amends the exemption to provide that the Act does not apply to a product, service, website, or application that provides exclusively nonuser-generated videogame systems incapable of producing highly realistic videos or images that a reasonable person could confuse with reality.

11) Makes other technical and clarifying changes to the CAIT Act.

12) 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**

Last session, the Legislature responded to some of these issues by passing the CAIT Act, SB 942 (Becker, Ch. 291, Stats. 2024), which is set to become operative, in part, later this year. The CAIT Act imposes obligations on "covered providers," persons that create, code, or otherwise produce a GenAI system that has over 1 million monthly visitors or users and is publicly accessible within the geographic boundaries of the state. It requires such providers to make an 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.

The CAIT Act also regulates. Covered providers are required to include a latent disclosure in AI-generated image, video, or audio content that is created by their GenAI system that is detectable using the above tool, and that is, to the extent technically feasible, permanent or extraordinarily difficult to remove. This latent disclosure must identify the provider, the tool, and the time and date of the content's creation or alteration. Covered providers are also required to provide users making such content with their system with the option to include a manifest disclosure that identifies it as AI-generated content.

Last year, AB 853 (Wicks, Chapter 674, Statutes of 2025) bolstered the CAIT Act by establishing similar transparency requirements on large online platforms, capture device manufacturers, and GenAI system hosting platforms.

This bill seeks to clean up the CAIT act and respond to some concerns raised by stakeholders. It includes a number of technical and clarifying amendments. Specifically, the bill removes the requirement for covered providers to provide the option of including a manifest disclosure in the relevant content. There were concerns that such perceptible labels would be easy to spoof and could lead to confusion for users.

This bill ensures that more minor AI alterations do not trigger the requirements of this bill. It also allows users to see personal information in provenance data if the user to whom the information pertains expressly consents to such information being included in the provenance, and is informed that such disclosures are irreversible once content is disseminated outside their control. This provides the option for users to include such information if, for instance, they are looking for attribution. This bill removes the user threshold for covered providers to ensure obligations for including disclosures also fall on those who download and modify open-source systems.

This bill also redefines "provenance data" to mean information about the origin of a piece of 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. The goal is to ensure the whole Act, as amended by AB 853, is aligned, and that large online platforms are deciphering such information.

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.

**SUPPORT:** (Verified 5/14/26)

ForensicVB LLC
Transparency Coalition.AI

**OPPOSITION:** (Verified 5/14/26)

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.

Prepared by: Christian Kurpiewski / P., D.T., & C.P. / (916) 651-1548
5/14/26 16:36:48

****** END ******