

SENATE PRIVACY, DIGITAL TECHNOLOGIES, AND CONSUMER PROTECTION COMMITTEE
Senator Christopher Cabaldon, Chair
2025-2026 Regular Session

AB 2713 (Wicks)
Version: April 21, 2026
Hearing Date: June 15, 2026
Fiscal: No
Urgency: No
CK

SUBJECT

California AI Transparency Act: system provenance data

DIGEST

This bill reworks certain provisions of the recently enacted California AI Transparency Act (CAIT) that apply to large online platforms.

EXECUTIVE SUMMARY

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, 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. One method that has been increasingly relied upon is embedding into content certain “provenance data,” which is data that is embedded into digital content, or that is included in the content’s metadata, for the purpose of verifying the digital content’s authenticity, origin, or history of modification.

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 content was created or altered by the covered provider’s system. Last year, AB 853 (Wicks, Ch. 674, Stats. 2025) bolstered the CAIT Act by establishing similar transparency requirements on large online platforms, capture device manufacturers, and GenAI system hosting platforms.

This bill reworks provisions added by AB 853 that apply to large online platforms. The bill is author-sponsored and supported by the Center for AI and Digital Policy. No timely opposition was received by the Committee.

PROPOSED CHANGES TO THE LAW

Existing law:

- 1) Establishes the California AI Transparency Act, which becomes operative, in part, on August 2, 2026, and requires certain “covered 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. (Bus. & Prof. Code § 22757 et seq.)
- 2) Requires a “covered provider,” a person that creates, codes, or otherwise produces a GenAI system that has over 1,000,000 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(a).)
- 6) Prohibits, starting January 1, 2027, a large online platform from, to the extent technically feasible, knowingly stripping any system provenance data or digital signature that is compliant with widely adopted specifications adopted by an established standards-setting body from content uploaded or distributed on the large online platform. (Bus. & Prof. Code § 22757.3.1(b).)
- 7) 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.)
- 8) 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. (Bus. & Prof. Code § 22757.5.)

This bill:

- 1) Reworks provisions of the CAIT Act that apply to large online platforms, making stylistic and technical changes.
- 2) Allows a large online platform to permit users to download content provenance data directly from the platform, in a format that cannot be easily integrated into unrelated content, rather than downloading content with provenance data attached as currently provided in the CAIT Act.
- 3) Prohibits a large online platform from knowingly stripping system provenance data or a digital signature from content a user *downloads* from the large online platform.

COMMENTS

1. Blurring reality: AI-generated content

GenAI can create new content, including text, images, code, or music, by learning from existing data. GenAI models can produce realistic and novel artifacts that resemble the data they were trained on, but do not copy it. The world has been in awe of the powers of this new technology, but the capabilities of these advanced systems lead to a blurring between reality and fiction. The National Institute for Standards and Technology frames the issue presented:

In recent years, the quality of digital content generated by artificial intelligence (AI) has advanced considerably across various modalities, including image, video, audio, and text. This surge in generative AI capability presents both opportunities and challenges – generative AI has facilitated creative expression and production, enabling artists, designers, and writers to create digital content at a much faster pace, but has also raised concerns regarding the authenticity and integrity of digital media, especially as it has become increasingly difficult to distinguish AI-generated content from human-generated content.¹

The relevant issue here is not about the growing incidence of synthetic content alone, but about the quickly blurring line between what is real and what is not. As one recent article phrased it, there is a “New Authenticity Crisis”:

OpenAI launched Sora 2 recently, an app that generates photorealistic video from text prompts. Within a week, Hollywood talent agencies called it exploitation. The Motion Picture Association objected to its copyright policies. Then came the flood of videos featuring copyrighted characters in unauthorized scenes. Despite being invite-only, the app hit one million downloads in five days. Faster than ChatGPT’s launch. Social media is now overwhelmed with synthetic clips so convincing that Zelda Williams, daughter of Robin Williams, publicly asked strangers to stop sending her AI-generated videos of her deceased father.

This is not an isolated technology story. It represents the acceleration of something larger happening across business and wider society.²

2. Cleaning up CAIT

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 1,000,000 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.

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

¹ NIST, *GenAI: Text 2026*, <https://ai-challenges.nist.gov/text-2026>. All internet citations are current as of June 6, 2026.

² Prashant Saxena, *Authenticity in the Age of AI* (December 22, 2025) California Management Review, <https://cmr.berkeley.edu/assets/documents/pdf/2025-12-authenticity-in-the-age-of-ai.pdf>.

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, Ch. 674, Stats. 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 rework provisions of the Act added by AB 853 that apply to large online platforms and respond to some concerns raised by stakeholders. It includes a number of technical and clarifying amendments, including extension of the disclosure requirements to provenance data not only embedded into or attached to content, but also provenance data otherwise associated with the content.

Currently, the law provides several options for platforms to disclose content provenance data to users. The bill amends one of those options and permits large online platforms to meet their obligations by allowing users to download provenance data attached to content rather than allowing them to download the content itself with the attached system provenance data. This must be provided in a format that cannot be easily embedded into, attached to, or associated with unrelated content. This comes in the wake of concerns about potential copyright issues and piracy of protected materials posted to these platforms.

In addition, the law currently prohibits a large online platform, to the extent technically feasible, from knowingly stripping any system provenance data or digital signature from content uploaded or distributed on the platform. This bill extends this prohibition to content downloaded from the platform.

According to the author:

New and emerging developments of generative AI (GenAI) tools have made it easier to create, edit, and doctor images, video, and audio. GenAI technologies can create and manipulate content to look realistic and convincing, which allow bad actors to create harmful content and spread disinformation. With the passage of AB 853, there would be more transparency of AI-generated content in the digital information ecosystem, and users would have more information to understand the source of content. AB 2713 would provide additional clarification within existing law in order that any requirements to share provenance information on large online platforms would have to be operationalized in a way that protects content creators and allow for large online platforms to provide meaningful information for users on their platforms.

The Center for AI and Digital Policy writes in support:

AB 2713 takes a lifecycle approach and requires that provenance data persists with synthetic content. Without this requirement, disclosures and labels collapse into formalities that fail at the point of greatest exposure, when content leaves controlled device environments and circulates in the broader information ecosystem. California's existing AI Transparency Act requires large online platforms to "provide a user interface that reliably indicates whether content was generated or substantially altered by a GenAI system" but its anti-stripping prohibition reaches only content uploaded or distributed on the platform. AB 2713 closes this gap by extending the anti-stripping prohibition to content downloaded from the platform.

SUPPORT

Center for AI and Digital Policy

OPPOSITION

None received

RELATED LEGISLATION

SB 1000 (Becker, 2026) makes a number of changes to the 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. SB 1000 is currently in the Assembly Privacy and Consumer Protection Committee.

SB 1050 (Ashby, 2026) makes it an unlawful practice for a person to create and cause to be published in an advertising medium an advertisement that includes a synthetic performer without a clear and conspicuous disclosure that the performer is synthetic. SB 1050 is currently pending referral in the Assembly.

SB 1146 (Gonzalez, 2026) requires, in part, an advertisement that includes the image, audio, or video of a digital replica or synthetic performer depicted as a health care provider that is generated or substantially altered using AI or other computer technology to promote the sale of a health-related consumer product or service to include a clear and conspicuous disclosure that the image, audio, or video, as applicable, of the person in the advertisement was generated or substantially altered by AI and that no human health care provider is depicted. SB 1146 is currently in the Assembly Privacy and Consumer Protection Committee.

AB 2713 (Wicks)
Page 7 of 7

AB 853 (Wicks, Ch. 674, Stats. 2025) *See Comment 2.*

SB 942 (Becker, Ch. 291, Stats. 2024) *See Comment 2.*

PRIOR VOTES:

Assembly Floor (Ayes 74, Noes 0)

Assembly Privacy and Consumer Protection Committee (Ayes 15, Noes 0)
