

Date of Hearing: June 23, 2026

ASSEMBLY COMMITTEE ON PRIVACY AND CONSUMER PROTECTION

Rebecca Bauer-Kahan, Chair

SB 1159 (Cabaldon) – As Amended June 9, 2026

PROPOSED AMENDMENTS

SENATE VOTE: 38-0

SUBJECT: Artificial intelligence: transparency and governance

SYNOPSIS

California law guarantees broad public access to government through public records laws, open meeting laws, and agency rulemaking procedures. These laws generally allow a “person” to request records, submit comments, appear at hearings, or otherwise participate in governmental decisionmaking. Over the past few years, the development of artificial intelligence (AI) – and especially generative AI – has facilitated public participation by helping individuals draft comments, prepare testimony, or request information from government agencies. These tools may be especially valuable for individuals who face barriers to participation, including people with disabilities. At the same time, AI can also be used to generate public engagement at a scale and speed far beyond ordinary human participation. Recent reports have raised concerns that AI-generated comments, objections, or records requests could be used to overwhelm agencies, misrepresent the views of real people, and distort the public record.

SB 1159 clarifies that for purposes of public engagement with government, “person” does not include artificial intelligence systems, automated agents, or robots. In doing so, the bill seeks to grant government agencies a tool to combat fraudulent participation. The bill is supported by a wide variety of local governments, and is opposed by Oakland Privacy and First Amendment Coalition on the grounds that it erects barriers to public participation without meaningfully addressing the issue of fraud.

Proposed Committee amendments, described in Comment #4, would narrow the bill by expressly preserving a natural person’s ability to use AI to facilitate their own engagement, making it unlawful to falsely represent that a natural person engaged with government, and making conforming changes to align with the AI Transparency Act. If passed by this Committee, the bill will next be referred to the Judiciary Committee.

EXISTING LAW:

- 1) Defines “person” to include any natural person, corporation, partnership, limited liability company, firm, or association for purposes of both the following:
 - a) The California Public Records Act. (Gov. Code § 7920.520.)
 - b) The Legislative Open Records Act. (Gov. Code § 9072.)
- 2) Pursuant to the California Public Record Act, requires each agency to, upon a request for a copy of records, within 10 days from receipt of the request, determine whether the request, in whole or in part, seeks copies of disclosable public records in the possession of the agency,

and to promptly notify the person making the request of the determination and the reasons therefor, except as specified. (Gov. Code § 7922.535.)

- 3) Pursuant to the California Legislative Records Act, requires that all requests to inspect any legislative record be made to the appropriate Rules Committee of each house of the Legislature or the Joint Rules Committee, that such committees shall be responsible for making all legislative records available for inspection, and that such legislative records shall be made available for inspection promptly and without unnecessary delay, except as specified. (Gov. Code § 9074.)
- 4) Pursuant to the Bagley-Keene Open Meeting Act, provides that it is the public policy of this state that public agencies exist to aid in the conduct of the people's business and the proceedings of public agencies be conducted openly so that the public may remain informed. (Gov. Code § 11120.)
- 5) Pursuant to the Administrative Procedures Act, requires state agencies to allow for public comment as follows:
 - a) If a public hearing is held, both oral and written statements, arguments, or contentions, shall be permitted.
 - b) If a public hearing is not scheduled, the state agency shall afford any interested person or his or her duly authorized representative, the opportunity to present statements, arguments or contentions in writing. (Gov. Code § 11342.575.)
- 6) Pursuant to the Ralph M. Brown act, provides that for local agencies, every agenda for regular meetings shall provide an opportunity for members of the public to directly address the legislative body on any item of interest to the public. (Gov. Code § 54954.3.)
- 7) For purposes of the California Environmental Quality Act:
 - a) Defines "person" to include any person, firm, association, organization, partnership, business, trust, corporation, limited liability company, company, district, county, city and county, city, town, the state, and any of the agencies and political subdivisions of those entities, and, to the extent permitted by federal law, the United States, or any of its agencies or political subdivisions. (Pub. Res. Code § 21066.)
 - b) Requires a lead agency to consider comments it receives on a draft environmental impact report, proposed negative declaration, or proposed mitigated negative declaration if those comments are received within the public review period. (Pub. Res. Code § 21091.)
- 8) For purposes of the California Coastal Act:
 - a) Defines "person" to mean any individual, organization, partnership, limited liability company, or other business association or corporation, including any utility, and any federal, state, local government, or special district or an agency thereof. (Pub. Res. Code § 30111.)
 - b) Provides that during the preparation, approval, certification, and amendment of any local coastal program, the public, as well as all affected governmental agencies, including

special districts, shall be provided maximum opportunities to participate. (Pub. Res. Code § 30503.)

- 9) Under the AI Transparency Act, beginning August 2, 2026, requires a creator of a GenAI system to additionally create and make freely available an AI detection tool that can be used to assess whether digital content was created or altered by a creator's GenAI system. (Bus. & Prof. Code § 22757.1 *et seq.*)
- 10) Defines "assistive technology" to mean any item, piece of equipment, or product system, whether acquired commercially, modified, or customized, that is used to increase, maintain, or improve functional capabilities of individuals with disabilities, and any service that directly assists an individual with a disability in the selection, acquisition, or use of the item, equipment, or product system. (Welf. & Inst. Code § 19461.)
- 11) Provides that it is unlawful for any person to use a bot to communicate or interact with another person in California online, with the intent to mislead the other person about its artificial identity for the purpose of knowingly deceiving the person about the content of the communication in order to incentivize a purchase or sale of goods or services in a commercial transaction or to influence a vote in an election. (Bus. & Prof. Code § 17941.)

THIS BILL:

- 1) In each of the following code sections, clarifies that "person," "interested person," "member of the public," and any other similar terms referring to those who may engage with governmental agencies do not include artificial intelligence systems, autonomous agents, or robots, whether physical or digital:
 - a) The California Public Records Act.
 - b) The Legislative Open Records Act.
 - c) The Bagley-Keene Open Meeting Act.
 - d) The Ralph M. Brown Act.
 - e) The California Environmental Quality Act.
 - f) The California Coastal Act.
- 2) In each of the code sections listed above:
 - a) Defines "artificial intelligence" to mean an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.
 - b) Provides that an agency may use an artificial intelligence detection tool that meets the criteria described in Section 22757.2 of the Business and Professions Code to determine if artificial intelligence is present.
- 3) Makes various findings and declarations.

COMMENTS:1) **Author's statement.** According to the author:

AI slop drowns out the voices of genuine human citizens trying to make their voices heard. SB 1159 puts an end to the intentional corruption of public engagement guarantees by those using AI agents to overwhelm public agencies, divert attention from real constituents, and strip public discourse of any meaning at all.

AI swarming has emerged as a real problem. Recently, the Los Angeles Times reported that at least 20,000 public comments generated by AI may have convinced Southern California's top air pollution authority to scrap a plan to phase out gas-powered appliances. The Times confirmed a public affairs consultant used a company that bills itself as an "AI-powered grassroots advocacy platform" to send the messages. The same playbook is being used to target the Bay Area air district, where a Chronicle investigation found that the residents who had comments submitted under their names never submitted anything. "This was forged," one San Pablo resident told the Chronicle. "I never wrote the letter." In the United Kingdom, an AI service called Objector.ai is generating similar concerns about swarming government with infinite objections ("AI-powered nimbyism could grind UK planning system to a halt, experts warn", The Guardian).

2) **Background.** *AI and GenAI.* The development of GenAI is creating exciting opportunities to grow California's economy and improve the lives of its residents. GenAI can generate compelling text, images and audio in an instant – but with novel technologies come novel safety concerns.

In brief, AI is the mimicking of human intelligence by artificial systems such as computers. AI uses algorithms – sets of rules – to transform inputs into outputs. Inputs and outputs can be anything a computer can process: numbers, text, audio, video, or movement. AI is not fundamentally different from other computer functions; its novelty lies in its application. Unlike traditional computer functions, AI can accomplish tasks that are normally performed by humans. AI that is trained on small, specific datasets to make recommendations and predictions is sometimes referred to as "predictive AI." This differentiates it from GenAI, which is trained on massive datasets in order to produce detailed text and images. When Netflix suggests a TV show to a viewer, that recommendation is produced by predictive AI that has been trained on the viewing habits of Netflix users. When ChatGPT generates text in clear, concise paragraphs, it uses GenAI that has been trained on the written contents of the internet.

Deepfakes and Disinformation. Image manipulation and video doctoring have existed for nearly as long as photography and recording equipment, but they have historically required great effort and talent. In the past few years, the rapid development of GenAI has drastically reduced those barriers to entry, allowing a vast quantity of convincing, but ultimately fake, content to be generated in an instant. The creation of imagery, video, and audio by GenAI has the potential to change the world by automating repetitive tasks and fostering creativity. When employed by bad actors, however, these capabilities have the potential to destroy lives and destabilize societies.

Scams. GenAI-powered speech and video is driving a new era in scamming. These AI tools are often trained on publicly available data, and as a result, the more data a target has online, the easier it is to develop a passable imitation of them or their loved ones. This is especially true of wealthy clients, whose public appearances, including speeches, are often widely available on the

internet.¹ For example, a complicated scam utilizing both deepfake video and false audio was recently performed in Hong Kong. A multinational company lost \$25.6 million after employees were fooled by deepfake technology, with one incident involving a digitally recreated version of its chief financial officer ordering money transfers in a video conference call. Everyone present on the video call, except the victim, was a fake representation of real people. The scammers applied deepfake technology to turn publicly available video and other footage into convincing versions of the meeting's participants.²

In December 2024, the FBI issued a public service announcement warning about the potential dangers posed by GenAI.³ Among the concerns highlighted was the technology's ability to significantly lower the barrier for producing counterfeit documents, such as fake IDs, passports, and other fraudulent government-issued identification, which could greatly facilitate identity theft. Additionally, GenAI enables the creation of entirely fictitious profiles on social media and dating platforms, which can be used to exploit individuals both financially and emotionally.

Elections. Deepfake technology is being used around the world to spread disinformation and propaganda. This has already been observed in Slovakia, where deepfake audio influenced an election in 2023:

Days before a pivotal election in Slovakia to determine who would lead the country, a damning audio recording spread online in which one of the top candidates seemingly boasted about how he'd rigged the election. And if that wasn't bad enough, his voice could be heard on another recording talking about raising the cost of beer. The recordings immediately went viral on social media, and the candidate, who is pro-NATO and aligned with Western interests, was defeated in September by an opponent who supported closer ties to Moscow and Russian President Vladimir Putin.⁴

Similar deepfakes surfaced in the United States ahead of the 2024 presidential election. In July 2024, Elon Musk shared a video featuring an AI-generated voice clone of then-Vice President Kamala Harris, in which the fabricated voice claimed she was a "diversity hire" due to being a woman of color and that she "did not know the first thing about running a country."⁵ Although Musk admitted two days after posting that the video was intended as satire, the potential impact of such content on political campaigns remains a serious concern. In 2023, the Legislature enacted a trio of bills aimed at addressing elections deepfakes: AB 2655 (Berman, Ch. 261, Stats. 2024), which imposes removal and disclosure obligations on large online platforms during the

¹ Emily Flitter and Stacy Cowley, "Voice Deepfakes Are Coming for Your Bank Balance", *New York Times* (Aug. 30, 2023), www.nytimes.com/2023/08/30/business/voice-deepfakes-bank-scams.html.

² Harvey Kong, "'Everyone looked real': multinational firm's Hong Kong office loses HK\$200 million after scammers stage deepfake video meeting," *South China Morning Post* (Feb. 4, 2024), www.scmp.com/news/hong-kong/law-and-crime/article/3250851/everyone-looked-real-multinational-firms-hong-kong-office-loses-hk200-million-after-scammers-stage.

³ Department of Justice Federal Bureau of Investigations, "Criminals Use Generative Artificial Intelligence to Facilitate Financial Fraud" (Dec. 3, 2024), <https://www.ic3.gov/PSA/2024/PSA241203>.

⁴ Curt Devine, Donie O'Sullivan, Sean Lyngass, "A fake recording of a candidate saying he'd rigged the election went viral. Experts say it's only the beginning," *CNN* (Feb. 1, 2024), www.cnn.com/2024/02/01/politics/election-deepfake-threats-invs/index.html.

⁵ Ken Bensinger, "Elon Musk Shares Manipulated Harris Video, in Seeming Violation of X's Policies", *The New York Times* (July 27, 2024), <https://www.nytimes.com/2024/07/27/us/politics/elon-musk-kamala-harris-deepfake.html>.

four months leading up to an election; AB 2839 (Pellerin, Ch. 262, Stats. 2024), which prohibits the distribution of election materials containing certain types of deceptively altered digital content; and AB 2355 (Carillo, Ch. 260, Stats. 2024), which requires AI-altered campaign materials to include clear disclosures. The first two bills, however, have been blocked from being enacted by legal proceedings.

Fraudulent government engagement. The author points to several recent incidents involving AI-enabled fraudulent engagement with government to justify this bill. In 2025, the South Coast Air Quality Management District was flooded with what appeared to be more than 20,000 AI-generated public comments opposing proposed air quality rules that would have phased out gas-powered appliances:

The opposition appeared overwhelming: Tens of thousands of emails poured into Southern California’s top air pollution authority as its board weighed a June proposal to phase out gas-powered appliances. But in reality, many of the messages that may have swayed the powerful regulatory agency to scrap the plan were generated by a platform that is powered by artificial intelligence.

[. . .]

When staffers at the air district reached out to a small sample of people to verify their comments, at least three said they had not written to the agency and were not aware of any such messages, records show.

But the email onslaught almost certainly influenced the board’s June decision, according to agency insiders, who noted that the number of public comments typically submitted on agenda items can be counted on one hand.⁶

In the United Kingdom, an AI service called Objector.ai is driving similar concerns about overwhelming AI-generated objections to new housing development:

The government’s plan to use artificial intelligence to accelerate planning for new homes may be about to hit an unexpected roadblock: AI-powered nimbyism.

A new service called Objector is offering “policy-backed objections in minutes” to people who are upset about planning applications near their homes.

It uses generative AI to scan planning applications and check for grounds for objection, ranking these as “high”, “medium” or “low” impact. It then automatically creates objection letters, AI-written speeches to deliver to the planning committees, and even AI-generated videos to “influence councillors”.

⁶ Hayley Smith, “Southern California air board rejected pollution rules after flood of comments from AI-powered platform,” *Los Angeles Times*, (Feb. 17, 2026), <https://www.latimes.com/environment/story/2026-02-17/ai-powered-campaign-may-have-killed-key-vote-on-air-quality>.

Kent residents Hannah and Paul George designed the system after estimating they spent hundreds of hours attempting to navigate the planning process when they opposed plans to convert a building near their home into a mosque.⁷

AI Transparency Act. Seeking to enact a comprehensive content provenance framework to ensure transparency around the rapid online proliferation of AI-generated content, the Legislature passed SB 942 (Becker, Stats. 2024, Ch. 291). The law applies to “covered providers” – developers of publicly-accessible GenAI systems with over one million monthly users – requiring them to provide a free, publicly accessible tool that allows users to detect whether audio, image, or video content was generated or altered by their systems. The bill permits visible disclosures and requires latent, machine-detectable provenance data to be embedded into content generated by these providers’ GenAI systems.

While SB 942 focused primarily on developers of large GenAI systems, AB 853 (Wicks, Stats. 2025, Ch. 674) imposed disclosure requirements on large online platforms and manufacturers of devices capable of capturing digital content (e.g. digital cameras.) Together, these bills effectively require GenAI developers to disclose that content generated by their systems is “fake” beginning in mid-2026, recording device manufacturers to disclose that content captured by their devices is “real” beginning in early 2028, and large online platforms to prominently disclose this provenance data to users beginning in early 2027.

3) **What this bill would do.** This bill would clarify that, for purposes of various forms of public engagement with government agencies, “person” does not include an AI system, autonomous agent, or robot. According to the author, the bill was introduced to address concerns that AI systems capable of facilitating fraudulent engagement at scale could soon overwhelm government processes:

AI systems can now autonomously engage in governmental processes at scales and speeds that far exceed human capacity, submitting infinite automated records requests, generating mass public comments, filing automated petitions, or using real-time video to impersonate participants on platforms like Zoom. This overwhelms agencies, diverts public resources, crowds out genuine human voices, and undermines deliberative processes.

Recognizing that agencies may find it difficult to distinguish between genuine and artificial participation, the bill explicitly allows agencies to use AI detection tools developed pursuant to the AI Transparency Act.

The bill is supported by a wide array of cities and counties. A coalition consisting of the Rural County Representatives of California, California State Association of Counties, Urban Counties of California, California Municipal Clerks Association, California Association of Councils of Governments, and League of California Cities writes the following in support:

Emerging AI and Large Language Model (LLM) tools have now made it easier than ever for bad actors to misrepresent public opinion in governance decisions. According to reporting from multiple outlets, in June 2025 a political consultant used AI tools to generate twenty

⁷ Aisha Down and Robert Booth, “AI-powered nimbyism could grind UK planning system to a halt, experts warn,” *The Guardian*, (Nov. 9, 2025), <https://www.theguardian.com/politics/2025/nov/09/ai-powered-nimbyism-could-grind-uk-planning-system-to-a-halt-experts-warn>.

thousand emails to the South Coast Air Quality Management District regarding a regulatory proposal before the board. The AI generated emails purported to be from members of the public in South Coast AQMD's jurisdiction, but during an audit by air district staffers many said that they were not aware of signing onto any such messages. In an interview with a trade publication, the political consultant behind the advocacy effort bragged about how the volume "left the South Coast AQMD staff reeling." This is just one example of an emerging trend that poses real concerns for the quality of civic discourse.

On the other hand, the bill is opposed by Oakland Privacy and First Amendment Coalition due to fears that the bill may inadvertently stifle public participation in government proceedings. Oakland Privacy writes in opposition:

There is potentially some privacy harm in the determination of what public comments have human origin and which do not, which it seems from the bill language is a determination that would be made by public officials without any clear criteria. As mentioned above, bulk email comments in the thousands, tens of thousands, and millions, can and are submitted by advocates on pressing matters. They may be quite legitimate, or as described above, they may be fraudulent whether through analog fraud or AI fraud. It is not entirely clear how an elected or appointed body would make that determination, and in order to do so accurately they would likely need to violate the privacy rights of the commenting parties. Under the Brown Act, for example, individuals are not compelled to disclose their names or addresses to make a public comment and can do so under an alias if they so desire...

The First Amendment Coalition concurs, writing the following:

It is not clear how an elected or appointed body would make a determination about whether AI was used and in order to do so accurately they would likely need to interrogate records requesters and members of the public, which raises serious chilling concerns . . . We are concerned that a municipality or other public agency can declare a person or entity is an AI and therefore it is unnecessary to fulfill a records request or list a public comment in the record.

4) **Committee amendments.** No government agencies currently treat AI, autonomous agents, or robots as people, and as a result, it is not entirely clear what effect this bill will have. To the extent SB 1159 does alter individuals' ability to request information and participate in government proceedings, it is important to distinguish between "good" uses of AI – for example, to facilitate participation by an individual who would have otherwise lacked the ability to engage due to a disability – and bad, fraudulent uses. To enable this distinction, the author has agreed to amend the bill in the following ways.

Amendment 1 adds the following provision to each section clarifying that individuals can use AI to facilitate their own engagement with government agencies, including through the use of assistive technologies, as that term is defined in Welf. & Inst. Code § 19461:

(a) This section does not prohibit a natural person from using artificial intelligence, autonomous agents, or robots to facilitate the person's own engagement with a governmental agency, including through the use of assistive technologies, provided the volume and frequency of the engagement are reasonably consistent with ordinary participation by a natural person.

(b) For purposes of this section, “assistive technology” means an item, piece of equipment, or product system, whether acquired commercially, modified, or customized, that is used to increase, maintain, or improve functional capabilities of individuals with disabilities and any service that directly assists an individual with a disability in the selection, acquisition, or use of the item, equipment, or product system.

Amendment 2 adds the following standalone provision clarifying that fraudulent use of AI is illegal with respect to all forms of government participation:

(a) It is unlawful for a person to knowingly use artificial intelligence to falsely represent that a natural person appeared before, submitted information to, or otherwise engaged with a governmental agency.

(b) For purposes of this section, “artificial intelligence” means an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.

Amendment 3 aligns terminology with SB 1000 (Becker), which seeks to amend the AI Transparency Act prior to that law going into effect this August:

(c) A state body may use an ~~artificial intelligence detection~~ *a disclosure verification tool* that meets the criteria described in Section 22757.2 of the Business and Professions Code to determine if artificial intelligence is present.

ARGUMENTS IN SUPPORT: A coalition of associations representing cities and counties, including CalCities and the California State Association of Counties, writes in support:

Automated engagement tools have the capacity to overwhelm government agencies, crowd out the voices of actual members of the public, and undermine the intent of California’s transparent governance laws. Clarifying how these laws apply in the context of emerging technologies, particularly distinguishing between authentic public input and activity generated at scale by artificial intelligence systems, represents an important step toward protecting meaningful public participation, open governance, and critical staff resources for local agencies. With thoughtful refinement, SB 1159 will help address these challenges in a way that is both effective and implementable as technology continues to evolve quickly.

ARGUMENTS IN OPPOSITION: First Amendment Coalition writes in opposition:

When it comes to public records requests, especially if one is filing them with multiple agencies to track how many agencies in California are doing a particular thing in order to assemble the necessary facts to advocate for policy change, a person or group may wish to use an AI tool to draft a request or to file the request, just as public agencies have been investing in AI tools to streamline their work. SB 1159 runs the risk of deterring important statewide research using public records by reporters, advocates, think tanks or academic researchers who may find value in an AI tool. SB 1159 is unnecessary especially as to the California Public Records Act, which already allows governmental agencies to deem requests over broad or too burdensome, and to work with requesters to target or narrow a request. The

problems potentially presented by AI, including fraudulent bulk emails and high volume public records requests, all have their analog equivalents and existing tools to address them.

Oakland Privacy writes in opposition:

The California Public Records Act already contains language to allow governmental agencies to defray requests they consider “over-broad” or “burdensome” and to relay back to those requesters a request to target or narrow that request. That language is a perfectly adequate solution to the problem the author is trying to solve and it already exists in statute. Public agencies in California already have the ability to label requests as over-broad and not to respond, just as city clerks for years have dealt with City Hall gadflies who come by every day asking for printouts or to speak on every single item in public comment. Not all problems are new, some have existed for time immemorial.

REGISTERED SUPPORT / OPPOSITION:

Support

Bay Area Air Quality Management District
 California Acupuncture Board
 California Association of Councils of Governments
 California Association of Recreation & Park Districts
 California Big City Mayors Coalition
 California Municipal Clerks Association (CMCA)
 California Police Chiefs Association
 California Special Districts Association
 California State Association of Counties (CSAC)
 California Yimby
 Center for Ai and Digital Policy (CAIDP)
 Cft – a Union of Educators & Classified Professionals, Aft, Afl-cio
 Chamber of Progress
 City of Bakersfield
 City of Beaumont
 City of Carlsbad
 City of Corona
 City of El Cerrito
 City of Merced
 City of Rancho Cucamonga
 City of Thousand Oaks
 Contra Costa County
 Dental Board of California
 Fresno County Board of Supervisors
 League of California Cities
 Marin County Council of Mayors and Councilmembers
 Riverside; City of
 Rural County Representatives of California (RCRC)
 Streets for All
 Urban Counties of California (UCC)

Opposition

First Amendment Coalition
Oakland Privacy

Analysis Prepared by: Slater Sharp / P. & C.P. / (916) 319-2200