

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

SB 1159 (Cabaldon)
Version: March 25, 2026
Hearing Date: April 6, 2026
Fiscal: No
Urgency: No
CK

SUBJECT

Artificial intelligence: transparency and governance.

DIGEST

This bill clarifies that, for purposes of specified transparency and governance laws, terms such as “person,” “interested person,” “member of the public,” and other similar terms do not include AI systems, autonomous agents, robots, or other nonhuman entities, whether physical or digital.

EXECUTIVE SUMMARY

California prides itself on open, transparent government that allows for vigorous public participation and oversight. This is borne out by the intricate network of transparency and governance laws that are intended to carry out the mandate found in the state’s Constitution, which states that “the people have the right of access to information concerning the conduct of the people’s business.” These laws include the California Public Records Act, the Bagley-Keene Open Meeting Act, the Ralph M. Brown Act, the Administrative Procedure Act, the California Environmental Quality Act, and the Political Reform Act of 1974.

These laws provide myriad methods to engage with and oversee the work of governmental bodies, including public comment periods and records requests. However, concerns have arisen that breakthroughs in agentic AI may subvert the principles underlying these laws by allowing any individual to click a button and overwhelm our current systems, flooding government offices with requests and creating a sea of public comments that threatens to drown out the voices of actual Californians.

This bill responds by making clear that various terms in these laws, such as “person” or “member of the public,” referring to those who may engage with governmental agencies under those laws, refer to natural persons and legally recognized entities capable of genuine participation in democratic governance, not AI systems that could be programmed to simulate participation at scales that would overwhelm governmental

processes. This bill is author-sponsored. It is supported by a broad coalition of government entities and advocacy groups. Oakland Privacy is in opposition. The bill passed out of the Senate Judiciary Committee on a 12 to 0 vote.

PROPOSED CHANGES TO THE LAW

Existing law:

- 1) Provides, pursuant to the California Constitution, that the people have the right of access to information concerning the conduct of the people's business, and, therefore, the meetings of public bodies and the writings of public officials and agencies are required to be open to public scrutiny. (Cal. const. art. I, § 3(b)(1).)
- 2) Governs the disclosure of information collected and maintained by public agencies pursuant to the California Public Records Act (CPRA). (Gov. Code §§ 7920.000 et seq.)
- 3) Establishes the Bagley-Keene Open Meeting Act (Bagley-Keene), which requires state bodies to conduct their business in open public meetings, except as provided, and establishes requirements and procedures for such meetings. (Gov. Code § 11120 et seq.)
- 4) Establishes the Ralph M. Brown Act (Brown Act), which secures public access to the meetings of public commissions, boards, councils, and agencies in the state. (Gov. Code §§ 54950 et seq.)
- 5) Establishes the California Environmental Quality Act (CEQA) that, generally, requires a public agency to prepare, or cause to be prepared, and to certify the completion of, an environmental impact report on a project that it proposes to carry out or approve that may have a significant effect on the environment or to adopt a negative declaration if it finds that the project will not have significant effects. Authorizes an action or proceeding to attack, review, set aside, void, or annul the acts or decisions of a public agency on the grounds of noncompliance with CEQA, as provided. (Pub. Res. Code § 21100 et seq.)
- 6) Establishes the Administrative Procedure Act (APA), which governs the adoption, amendment, or repeal of regulations by state agencies and for the review of those regulatory actions by the Office of Administrative Law. (Gov. Code §§ 11340 et seq.)
- 7) Establishes the Political Reform Act of 1974 (PRA), which creates the Fair Political Practices Commission and makes it responsible for the impartial, effective administration and implementation of the PRA. (Gov. Code §§ 81000 et seq.)

This bill:

- 1) Provides that for the purposes of the CPRA, Bagley-Keene, Brown Act, CEQA, APA, and the PRA, the terms “person,” “interested person,” “participant,” “member of the public,” as applicable, and any other similar terms under each act referring to those who may engage with governmental agencies, do not include AI systems, autonomous agents, robots, or other nonhuman entities, whether physical or digital.
- 2) 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.
- 3) States relevant findings and makes clear that the changes made hereby are not to be construed to imply that the terms “person,” “interested person,” “participant,” “member of the public,” and any other similar terms as used in any law other than those laws referenced above include AI systems, autonomous agents, robots, or other nonhuman entities, whether physical or digital.

COMMENTS

1. An AI agent goes to Washington

Democratic governments have long guaranteed citizens transparency, access to the government and its decision-making processes, and the ability to scrutinize power. Public records laws let anyone demand government documents. Public comment periods on proposed legislation, ordinances, and other regulations invite ordinary people to weigh in before laws and rules take effect. These mechanisms were designed for a world where making requests and writing comments took time, effort, and human attention. That world is changing as society enters the “Agentic AI age.”

Modern AI agents are software systems that can be assigned tasks, break them into steps, and carry out those steps autonomously. These agents can be provided a series of tools and Internet access to accomplish their assignments, which might include browsing websites, filling out forms, generating text and submitting it, and interacting with other agents and people, all without pausing for human review:

Rewind a few years, and large language models and generative artificial intelligence were barely on the public radar, let alone a catalyst for changing how we work and perform everyday tasks.

Today, attention has shifted to the next evolution of generative AI: AI agents or agentic AI, a new breed of AI systems that are semi- or fully autonomous and thus able to perceive, reason, and act on their own.

Different from the now familiar chatbots that field questions and solve problems, this emerging class of AI integrates with other software systems to complete tasks independently or with minimal human supervision.

“The agentic AI age is already here. We have agents deployed at scale in the economy to perform all kinds of tasks,” said Sinan Aral, a professor of management, IT, and marketing at MIT Sloan.

Nvidia CEO Jensen Huang, in his keynote address at the 2025 Consumer Electronics Show, said that enterprise AI agents would create a “multi-trillion-dollar opportunity” for many industries, from medicine to software engineering.¹

Despite these opportunities and the potential beneficial uses throughout society, concerns have arisen that AI agents may upend the traditional ways in which the public engages with their government and change the arithmetic of civic participation in a fundamental way.

For instance, a single motivated person might file dozens of public records requests in a year and a well-resourced advocacy group might file hundreds. Generally, each request legally obligates a government entity to search its files, redact sensitive material, and respond within a statutory deadline. This work can take staff hours or days. An AI agent, given a target and a set of topics, could file thousands of requests in an afternoon. Because most public records laws have no mechanism to distinguish a human requester from an automated one, agencies could be legally required to treat each submission the same. This could bury staff in a sea of requests, and just as importantly, could interfere with legitimate requests, such as from journalists and researchers. The tool designed to create transparency could be weaponized to produce paralysis.

Public comment periods face a similar vulnerability. For instance, many governmental entities propose new rules or regulations and typically open a public-comment period during which any member of the public can submit a comment that the agency is legally required to read and consider before finalizing the rule. If AI agents can generate exponentially more than any human can, the legal and administrative burden of processing them would be immense, and the signal-to-noise ratio for genuine public sentiment collapses entirely.

These concerns are not theoretical, as revealed by one recent incident:

The opposition appeared overwhelming: Tens of thousands of emails poured into Southern California’s top air pollution authority as its board

¹ Beth Stackpole, *Agentic AI, explained* (February 18, 2026) MIT Management Sloan School, <https://mitsloan.mit.edu/ideas-made-to-matter/agentic-ai-explained>. All internet citations are current as of March 28, 2026.

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.

Public records requests reviewed by The Times and corroborated by staff members at the South Coast Air Quality Management District confirm that more than 20,000 public comments submitted in opposition to last year's proposal were generated by a Washington, D.C.-based company called CiviClick, which bills itself as "the first and best AI-powered grassroots advocacy platform."

A Southern California-based public affairs consultant, Matt Klink, has taken credit for using CiviClick to wage the opposition campaign, including in a sponsored article on the website Campaigns and Elections. The campaign "left the staff of the Southern California Air Quality Management District (SCAQMD) reeling," the article says....

For years, companies have employed bots or orchestrated fake "astroturf" campaigns to create the appearance of grassroots opinion on an issue, but the introduction of AI technology could make it even harder for elected officials to engage in earnest with the public.

"What we're seeing with AI is absolutely the next step in digital astroturfing," said Samuel Woolley, a researcher at the University of Pittsburgh who studies disinformation and the use of emerging technology in politics.²

However, the problem is not new, as AI bots flooded the system almost a decade ago when the federal government was dealing with net neutrality regulations:

In 2017, during a Federal Communications Commission (FCC) public comment period, bots flooded the agency with more than a million comments from fake constituents calling for the repeal of net neutrality rules. Regulators uncovered this covert attempt to subvert the policy process, but only because hundreds of thousands of the comments were uncannily similar. In the six years since, artificial intelligence (AI) tools have only grown more sophisticated, and similar efforts to deceive policymakers promise to become increasingly difficult to detect and prevent. This troubling risk and its implications for participatory government call for urgent action.

² Hayley Smith, *Southern California air board rejected pollution rules after flood of comments from AI-powered platform* (February 17, 2026) Los Angeles Times, <https://www.latimes.com/environment/story/2026-02-17/ai-powered-campaign-may-have-killed-key-vote-on-air-quality>.

The practice of open and responsive government is integral to the American tradition. Underlying it is the principle that the public should have ample access to information and opportunity to weigh in on important policy decisions. New technologies, from the invention of the telegraph to the explosion of social media, have produced new pathways for public democratic engagement. Yet these same technologies often create new tools that can be used to misrepresent public opinion and distort policymaking. AI is no exception.

As with prior technological revolutions, developments in AI hold promise: in time, they could enhance government's ability to understand what the public wants and to help citizens participate in policy decisions. In particular, automated summation and analysis could be deployed to augment officials' capacity to digest public feedback. But these potential benefits can only be realized alongside guardrails to ensure that AI systems are accurate and fit for the purposes at hand, to prevent system biases from reflecting sociohistorical inequities and diminishing the voices of specific communities, and to mitigate the effects when such outcomes occur.³

Agentic AI has only exacerbated the problem, which is ultimately structural. These government systems were not designed with volume attacks in mind because volume attacks were not possible at scale by private actors. Authentication requirements are minimal in many of these situations and the legal frameworks themselves, enshrined in statutes like the Administrative Procedure Act and the Public Records Act, assume that the challenge of participation is getting people to engage, not filtering out synthetic engagement. As stated in the bill's findings: "California's transparency and governance laws . . . presuppose participants who possess consciousness, moral agency, deliberative judgment, and membership in the political community. AI systems, regardless of their sophistication, lack these essential attributes of personhood."

Governments are beginning to grapple with this, but slowly. Some agencies are exploring verification requirements, CAPTCHA systems, or flagging submissions that arrive in statistically improbable patterns. For now, the mismatch between the speed of AI agents and the pace of bureaucratic and legislative reform leaves a significant gap. The machinery of public participation was built for humans. It has not yet been rebuilt for a world where humans are optional.

³ Mekela Panditharatne & Daniel Weiner, *Artificial Intelligence, Participatory Democracy, and Responsive Government* (November 3, 2023) Brennan Center for Justice, <https://www.brennancenter.org/our-work/research-reports/artificial-intelligence-participatory-democracy-and-responsive-government>.

2. Robots need not apply

This bill addresses these issues simply by making clear that for California's governance and transparency laws, namely the California Public Records Act, the Bagley-Keene Open Meeting Act, the Ralph M. Brown Act, the Political Reform Act of 1974, the Administrative Procedure Act, and CEQA, terms such as "person," "interested person," "participant," "member of the public," and other similar terms, as applicable, do not include AI systems, autonomous agents, robots, or other nonhuman entities, whether physical or digital. The bill includes a series of findings and declarations supporting this approach. To avoid any misinterpretation, the bill makes clear that these changes are not to be construed to imply that the use of such terms in any other law include AI systems, autonomous agents, robots, or other nonhuman entities.

According to the author:

SB 1159 advances a simple but urgent principle: California's transparency and governance laws were written for people, and they should work for people. As artificial intelligence systems grow more capable of mimicking human participation at scales no human could match, we must be clear that automated systems are not the "persons" our democratic institutions were designed to serve.

California's Constitution guarantees the people the right to access information about the conduct of their government's business. But that guarantee cuts both ways; it also depends on the government remaining capable of meaningfully responding to the public. When an agency is flooded with thousands or millions of AI-generated comments, records requests, or petitions, it cannot do that. AI-generated submissions divert limited public resources, drown out genuine human voices, and compromise the deliberative processes our laws were designed to protect.

We have already seen this happen. In 2025, the South Coast Air Quality Management District received more than 20,000 AI-generated public comments opposing air quality rules, comments that agency insiders say almost certainly influenced the Board's decision to scrap those rules. Some of the people whose names were used were unaware that anyone had submitted messages in their name. This is not public participation. It is digital astroturfing at a scale that our existing laws were never designed to address. SB 1159 ensures those voices are not drowned out by a flood of machine-generated noise and that California's transparency laws continue to serve the people for whom they were written.

3. Stakeholder positions

A coalition of groups representing local governments, including the League of California Cities and the California State Association of Counties, explain the need for the bill:

California's open meeting and open governance laws are intended to not only allow the public to observe deliberations about issues that will affect their communities, but to participate as well. Local agencies accept public comment on matters both on the agenda and off. Technology has been a powerful tool to make public participation more accessible than ever, as agendas and the time and location of public meetings are posted online. Those unable to comment or observe in person are able to follow along online or submit written comments via email or web portal.

However, 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. . . . 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.

CalNonprofits writes in support:

AI systems can now autonomously engage in governmental processes at scales and speeds that far exceed human capacity. These systems can submit infinite automated public records requests, generate mass public comments on proposed regulations, file automated petitions for rulemaking, or otherwise flood governmental agencies with interactions that simulate human participation. . . .

As technology advances, government bodies need to update and adapt to how they view civic engagement. SB 1159 preserves the integrity and functionality of California's democratic institutions by preventing automated systems from displacing genuine human participation, protecting public resources from being consumed by responding to machine-generated requests, and ensuring that governmental decisions remain responsive to the people of California.

Oakland Privacy writes in opposition:

In short, we don't think that the advent of artificial intelligence renders California's governmental transparency legal structure moot and in need of serious revision. The problems potentially presented by AI, including fraudulent bulk emails, lengthy Zoom public comment sessions, and high volume public records requests all have their analog equivalents and existing tools to address them. They are not, like most things, perfect, and the author knows as well as anyone that it can be exhausting to be a public servant, but they have on the whole served California well for the past 70+ years.

We do not think Californians should have to worry that the newest response to speech that the government doesn't like is "you're not human, you're an AI".

SUPPORT

California Special Districts Association
California State Association of Counties
California YIMBY
CalNonprofits
CFT - a Union of Educators & Classified Professionals, AFT, AFL-CIO
Chamber of Progress
City of Placentia
City of San Jose
City of Stanton
City of Westminster
County of Napa
County of Sacramento
County of Yolo
League of California Cities
Rural County Representatives of California
Streets for All
TechEquity Action
Urban Counties of California

OPPOSITION

Oakland Privacy
1 Individual

RELATED LEGISLATION

AB 410 (Wilson, 2025) requires that bots disclose their identity before autonomously interacting with a person, respond truthfully to any query about their identity, and otherwise refrain from misrepresenting themselves as human. AB 410 is currently in the Senate Appropriations Committee.

SB 1001 (Hertzberg, Ch. 892, Stats. 2018) prohibits a person from using a bot to communicate or interact with another person in California online with 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. “Bot” is defined as an automated online account where all or substantially all of the actions or posts of that account are not the result of a person.

PRIOR VOTES:

Senate Judiciary Committee (Ayes 12, Noes 0)
