Date of Hearing: July 1, 2026

ASSEMBLY COMMITTEE ON PRIVACY AND CONSUMER PROTECTION
Rebecca Bauer-Kahan, Chair
SB 813 (McNerney) – As Amended June 16, 2026

PROPOSED AMENDMENTS

SENATE VOTE: 31-7

SUBJECT: California Artificial Intelligence Standards and Safety Commission: artificial intelligence safety standards

SYNOPSIS

"Audits allow us to make informed choices about the kind of world in which we want to live, and they move important questions about our lives and collective values into public dialogue… As AI systems increasingly shape whole swaths of our lives and livelihoods, audits are an ever more critical source of that information. The next step is up to all of us."

- *Auditing AI*[1]

Artificial intelligence systems are increasingly being deployed across high-impact sectors in California. While several prior or pending bills have proposed independent auditing requirements, no such requirement has yet been enacted in California. This gap is becoming more significant as industry, federal policymakers, and other states increasingly support requirements for independent evaluation of advanced AI systems.

SB 813 would establish the California Artificial Intelligence Standards and Safety Commission within the Business and Consumer Services Agency, and require the commission to develop and publish standards, procedures, and criteria for the assessment of AI systems and models by AI auditors. The commission would create both minimum compliance standards, focused on whether an AI system or model meets applicable California legal requirements, and advanced safety standards, focused on whether a system substantially exceeds minimum requirements and reflects emerging best practices. The bill would also require the commission to develop procedures for suspending or terminating AI auditor registration, and criteria for designating qualified auditors as independent verification organizations (IVOs) – effectively super-auditors capable of assessing whether a system or model goes above and beyond minimum compliance requirements.

This bill is sponsored by Fathom and supported by a variety of AI safety advocacy organizations. It is opposed by a variety of tech trade associations. Committee amendments, described in comment #4, would make various minor substantive and technical changes.

EXISTING LAW:

1) 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

[1] The Marquand House Collective, "Auditing AI," *MIT Press,* 2026.

receives how to generate outputs that can influence physical or virtual environments. (Bus. & Prof. § 22757.1.)

2) Establishes the AI Transparency Act and, beginning August 2, requires a person that creates, codes, or otherwise produces a generative artificial intelligence system to include a latent disclosure conveying specified information in image, video, or audio content generated by the system. (Bus. & Prof. Code § 22757 *et seq*.)

3) Establishes the Transparency in Frontier Artificial Intelligence Act and requires a large AI frontier developer, as defined, to write, implement, comply with, and clearly and conspicuously publish on its internet website a frontier AI framework that applies to the developer's frontier models and describes how the developer approaches specified high-risk situations. (Bus. & Prof. Code § 22757.10 *et seq*.)

4) Establishes in state government the Business, Consumer Services, and Housing Agency, consisting of specified departments and agencies, including the Department of Consumer Affairs. (Gov. Code § 12804.)

5) Establishes the California Board of Accountancy in the Department of Consumer Affairs, as specified. (Bus. & Prof. Code § 5000.)

THIS BILL:

1) Defines the following terms:

   a) "Advanced safety standard" means a standard published by the agency that establishes criteria for determining whether an AI system or model substantially exceeds minimum compliance requirements, incorporates emerging best practices, and achieves demonstrable outcomes that promote the safety, efficacy, reliability, and security of the system or model.

   b) "Agency" means the Business and Consumer Services Agency.

   c) "AI auditor" means a person, partnership, academic institution, nonprofit, or corporation that conducts a covered audit on behalf of a third party.

   d) Except as provided below, "commission" means the California Artificial Intelligence Standards and Safety Commission established pursuant to this bill.

      i. If AB 1709 (Lowenthal) of the 2025–26 Regular Session is enacted and establishes an e-Safety Advisory Commission, "commission" means the e-Safety Advisory Commission established pursuant to AB 1709 of the 2025–26 Regular Session.

   e) "Covered audit" means either of the following:

      i. An audit of an AI system or model conducted to verify whether the AI system or model complies with applicable state law.

      ii. An assessment of an AI system or model conducted in accordance with an advanced safety standard.

f) "Independent verification organization" or "IVO" means an AI auditor that is designated by the agency as having demonstrated expertise in assessing whether AI systems or models substantially exceed minimum compliance requirements, incorporate emerging best practices, and achieve demonstrable outcomes that promote their safety, efficacy, reliability, and security.

g) "Minimum compliance standard" means a standard published by the agency or created by statute that establishes criteria for determining whether an AI system or model meets minimum safety, efficacy, reliability, and security requirements necessary to comply with applicable state law.

2) Requires that on or before July 1, 2027, the agency shall establish the California Artificial Intelligence Standards and Safety Commission that consists of the following members:

a) A representative, appointed by the Governor, from a business that develops or deploys AI systems or models and employs more than 100 workers.

b) A representative, appointed by the Governor, from a business that develops or deploys AI systems or models and employs fewer than 100 workers.

c) An expert from academia, appointed by the Governor, with experience researching AI safety.

d) An individual, appointed by the Governor, representing the interests of civil society, including, but not limited to, nongovernmental organizations, public policy institutes, or consumer advocacy organizations.

e) An individual, appointed by the Governor, representing the legal and economic interests of workers in relation to AI systems and models.

f) The President of the California Board of Accountancy or that person's designee.

g) The Attorney General or that person's designee.

3) Provides that commission members shall serve four-year terms and may be reappointed.

4) Requires the commission to develop standards for the assessment of AI systems and models by AI auditors by doing all of the following:

a) Evaluate existing standards, frameworks, guidelines, and best practices developed or published by government agencies, standards-setting organizations, AI auditors, private entities that develop or deploy AI systems or models, or other independent experts with relevant expertise.

b) Establish a "minimum compliance standard" and an "advanced safety standard." In developing advanced safety standards, requires the commission to incorporate criteria for assessing all of the following with respect to an AI system or model in operation:

i. The safety, efficacy, reliability, security, and robustness of the AI system or model.

ii. Reasonably foreseeable risks arising from the development, deployment, operation, use, or misuse of the AI system or model.

iii. The adequacy of any measures adopted to prevent, mitigate, and respond to those risks.

iv. Whether the AI system or model has achieved and maintained a level of risk that is appropriate in light of the AI system's intended purpose, foreseeable uses and misuses, affected populations, and the severity and likelihood of potential harms.

v. Whether the AI system or model should be subject to ongoing monitoring, periodic reassessment, or other post-deployment review.

vi. Whether a developer, deployer, or other responsible entity has implemented and adhered to appropriate internal safety standards and protocols.

c) Prioritize, at a minimum, developing standards for AI systems, models, and use cases that are subject to state law.

d) Create a process for independent experts to submit advanced safety standards for evaluation by the commission pursuant to this chapter.

5) Requires the commission to develop procedures for determining whether to suspend or terminate the registration of an AI auditor by considering all of the following:

a) Failures to adhere to appropriate standards, including standards developed pursuant to this chapter.

b) Material misrepresentations in audit reports or required disclosures.

c) Conflicts of interest that impair independence.

d) Failure to maintain adequate documentation.

e) Conduct that reasonably calls into question the integrity, objectivity, or competence of the AI auditor.

6) Requires the commission to develop criteria for determining whether an AI auditor qualifies as an independent verification organization by considering whether the IVO is capable of all of the following with respect to an AI system or model in operation:

a) Assessing the safety, efficacy, reliability, security, and robustness of the AI system or model.

b) Identifying reasonably foreseeable risks arising from the development, deployment, operation, use, or misuse of the AI system or model.

c) Verifying that the AI system or model has achieved and will maintain a level of risk that is appropriate in light of the AI system's or model's intended purpose,

foreseeable uses and misuses, affected populations, and the severity and likelihood of potential harms.

d) Determining whether the AI system or model should be subject to ongoing monitoring, periodic reassessment, or other post-deployment review.

7) Requires the commission to publish any standards, procedures, and criteria developed pursuant to this chapter in a publicly accessible format on the agency's internet website, and to identify all of the following to the extent applicable:

a) Any state law to which the standard is intended to apply.

b) Whether the standard is a minimum compliance standard or an advanced safety standard.

c) The category, use case, or type of AI system or model to which the standard is intended to apply.

d) The date upon which the standard was adopted and the date of the most recent revision to the standard.

e) Any additional information necessary to facilitate the interpretation and application of the standard.

8) Requires the commission to publish a statement on the agency's internet website prominently disclosing that the publication of a standard pursuant to this subdivision does not constitute recommendation or endorsement by the state of any AI system or model.

9) Requires the commission to review standards, procedures, and criteria published pursuant to this chapter at least annually and revise them as appropriate to reflect changes in applicable state law, technological developments, widely recognized industry standards, and emerging best practices.

10) In carrying out its duties, permits the commission to consult and coordinate with any of the following:

a) Certified public accountants and public accountants in good standing as certified by the California Board of Accountancy.

b) AI auditors that are not certified public accountants or public accountants.

c) Independent verification organizations.

d) Academic institutions.

e) Private entities that develop or deploy AI systems or models.

f) Consumer protection, labor, and civil society organizations.

g) Relevant federal, state, and local agencies.

h) National and international standards-setting organizations.

i) Any other relevant stakeholders, as determined by the commission.

11) Requires the commission to convene working groups to solicit stakeholder input in the development and revision of standards, procedures, and criteria pursuant to this chapter.

12) Permits the commission to establish voluntary pilot programs to evaluate the effectiveness of standards developed pursuant to this chapter.

13) Provides that the bill does not do either of the following:

a) Establish liability solely for failure to comply with a standard developed or published pursuant to this chapter.

b) Constitute recommendation or endorsement by the state of any AI system or model.

14) Provides that the bill shall become operative only if Assembly Bill 1405 of the 2025–26 Regular Session is enacted.

15) Makes various findings and declarations.

COMMENTS:

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

California is a world leader in AI development, so it is incumbent on our state to ensure that the use of artificial intelligence is safe and beneficial. To do so, it is imperative that we establish strong yet workable standards — standards created by independent, third-party experts and academics who can nimbly adapt to evolving technology.

SB 813 is an innovative and pragmatic approach to ensuring that artificial intelligence is developed responsibly. With the public-private governance concept, we can advance high-level standards to improve consumer awareness and safety.

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.

AI auditing. Auditing may evoke visions of the IRS pursuing individuals or corporations for fraudulent tax practices, but audits are commonplace in numerous fields. Even the idea of auditing an algorithm itself is not new. In the 1960s, American Airlines began using a rudimentary computer program to aid in the booking of flights, which streamlined the process for travel agents and prevented double bookings. American Airlines extended this program to other airlines in the 1980s; however, an audit of this program discovered that the algorithm ensured that American Airlines flights would show up as the top result in any search, regardless of whether they were longer or more expensive. During a US Congressional investigation into this business practice, the President of American Airlines did not deny that the algorithm prioritized their flights. Instead, he doubled down, stating that "the preferential display of our flights, and the corresponding increase in our market share, is the competitive raison d'etre for having created the system in the first place."[2] Although regulations were ultimately put in place to address this algorithmic bias, there remains little to no legislative oversight for algorithms deployed in other sectors.

Audits of AI systems have often been in response to investigative journalism or public outcry. In 2016, ProPublica reported that COMPAS, a commercially available algorithm used to predict recidivism rates among criminals, predicted that Black individuals had a 77% higher likelihood of recidivism than white individuals, even when controlling for prior crimes, age, gender and actual recidivism. This finding suggested that race alone significantly influenced the algorithm's output and led to a reevaluation of the use of this algorithm.[3] In 2020, Twitter introduced an AI-powered image cropping tool designed to center photos in order to highlight faces or other interesting elements. However, users quickly observed that the tool disproportionately favored white faces over Black faces and women over men. This was later verified by Twitter.[4]

Despite these success stories, most AI systems remain "black boxes" whose design elements and efficacies are held as trade secrets, making it difficult for the public to evaluate whether these systems are being used legally and ethically. Testing these systems becomes challenging when only customer-facing outputs are accessible, especially as the number of variables involved increases: it is easy to verify if a calculator is adding two numbers correctly, but when an algorithm is used to assess college admissions, parsing the relative contributions of factors such as essays, grades, and extracurricular activities becomes nearly impossible.

In 2024, Anthropic, a major developer of large language models and frequent advocate for safety-oriented AI development, argued that independent third-party testing and auditing should become a central feature of frontier AI governance, but that such a regime should be developed carefully through iteration among industry, government, and academia:

> Developing a third-party testing regime for the AI systems of today seems to give us one of the best tools to manage the challenges of AI today, while also providing infrastructure we can use for the systems of the future. We expect that ultimately some form of third-party

---

[2] Christian Sandvig et al, "Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms," *Data and Discrimination: Converting Critical Concerns into Productive Inquiry* (May 22, 2014), accessed at https://kevinhamilton.org/.

[3] Jeff Larson et al., "How We Analyzed the COMPAS Recidivism Algorithm," *ProPublica* (May 23, 2016), access at https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm.

[4] Rumman Chowdhury, "Sharing learnings about our image cropping algorithm", *X Engineering* (May 19, 2021), accessed at https://blog.x.com/engineering/en_us/topics/insights/2021/sharing-learnings-about-our-image-cropping-algorithm.

> testing will be a legal requirement for widely deploying AI models, but designing this regime and figuring out exactly what standards AI systems should be assessed against is something we'll need to iterate on in the coming years - it's not obvious what would be appropriate or effective today, and the way to learn that is to prototype such a regime and generate evidence about it.[5]

Anthropic's position has since shifted toward calling for mandatory, binding pre-deployment testing and auditing for frontier models. In an essay published earlier this month, Anthropic CEO Dario Amodei argued these previously abstract risks are now "clearly here":

> It is time to go beyond transparency to more serious and binding regulation of AI. I believe the best analogy, at least at the current stage of the exponential, is to cars, airplanes, or drugs—powerful technologies essential to the modern economy, but capable of killing large numbers of people if designed or operated poorly. I therefore believe we should model AI regulation on agencies like the Federal Aviation Administration (FAA). **Frontier AI models, like airplanes, should be required to go through technical testing and auditing, and their release should be blocked or reversed as a threat to public safety if they do not meet high standards of safety.**
>
> . . .
>
> Models above a threshold of compute should undergo mandatory testing by a qualified third party for their level of risk in four specific areas: cybersecurity, biological weapons, loss of control of AI systems, and automated R&D that could accelerate these other risks.[6]

Alongside this essay, Anthropic released an "advanced AI framework" that included a requirement for developers of advanced AI systems to partner with independent evaluators. The framework described several key components for crafting a successful system of independent evaluation:

> **Require independent evaluation.** Within six months of the enactment of the regulation, require Covered Developers to regularly engage at least one qualified independent evaluator to examine the relevant covered models, gather information about safeguards, and provide their own assessment of Enumerated Risks.
>
> - The evaluator should be independent in the sense of not having a financial interest in the developer and being free of major conflicts of interest with respect to the individuals involved in conducting the review.
>
> - The evaluator should receive access to an unredacted version of the developer's most recent risk report and system cards, access to the developer's most capable models, and the opportunity to ask and receive reasonable responses to relevant questions about

---

[5] Anthropic, "Third-party testing as a key ingredient of AI policy," (Mar. 25, 2024), https://www.anthropic.com/news/third-party-testing

[6] Dario Amodei, "Policy on the AI exponential," (Jun. 2026), https://darioamodei.com/post/policy-on-the-ai-exponential

models, likelihood of catastrophic risks, and safeguards.

- The evaluator should publish a review of the developer's most recent risk report, addressing the adequacy of the report's information; its analytical rigor; the appropriateness and materiality of any redactions made to the public version; and whether the evaluator disagrees with any of the report's key claims, especially its overall assessment of the level of risk for each Enumerated Risk category.

- Evaluators should be bound by obligations not to copy, retain, or disclose confidential information (e.g., confidential intellectual property, matters of national security, or proprietary details such as model architecture, cost, or size), but beyond that should generally not be restricted in what they can publish, including concerns about the risk report or the developer's conduct in connection with the review process.

Independent evaluator ecosystem. Take steps to grow an independent evaluation ecosystem, including developing and publishing standards for evaluators; exploring a licensing system to qualify evaluators; providing government funding or arranging pooled funding so that evaluators can do their work while remaining financially independent of any given developer; and providing resources and funding for nascent organizations seeking to become evaluators.

Independent evaluation can reveal risks that would otherwise remain hidden from the public, regulators, and even downstream users – but only if auditors have the expertise, independence, access, and standards necessary to conduct meaningful reviews.

California's AI auditing bills. Over the past two years, as California has begun regulating the development and deployment of AI systems, independent auditing has repeatedly emerged as a proposed mechanism for translating legal requirements into practical, verifiable oversight:

- SB 53 (Wiener; Ch. 138, Stats. 2025), which established the "Transparency in Frontier Artificial Intelligence Act," requires developers of the most advanced, costly artificial intelligence systems to implement and disclose protocols for mitigating risks of catastrophic harms due to their models. Version 94 of SB 53 would have required large developers to annually retain an independent third-party auditor to produce a report assessing developer's compliance with their own safety and security protocol. Auditors would have been granted access to any materials reasonably necessary to perform the required assessment, and the results of the audit would have been transmitted to the Attorney General.

- AB 1018 (Bauer-Kahan, 2025), currently pending on the Senate Floor, would require developers of automated decision systems designed or used to make or facilitate "consequential decisions" to contract with independent third-party auditors to assess their compliance with the bill. Auditors would similarly be provided access to any available information that is reasonably necessary for the auditor to comprehensively assess compliance, and a high-level summary of the auditor's feedback would be publicly available at no cost to users of the developer's internet website.

- SB 503 (Weber-Pierson, 2025), currently pending on the Assembly floor, would require developers and health care deployers to make reasonable efforts to identify AI systems that

are known or have a reasonably foreseeable risk of biased impacts on health programs or activities. Version 94 of this bill would have required developers to submit their systems to independent third-party auditors to assess compliance, and to make a high-level summary of the audit results publicly available at no cost on the developer's internet website.

- AB 1064 (Bauer-Kahan, 2025) would have prohibited making a companion chatbot available to children if it is foreseeably capable of specified harmful behaviors, including encouraging the child to engage in self-harm, suicidal ideation, or violence, or engaging in sexually explicit interactions with the child. Version 97 of this bill would have required developers of covered products to regularly engage with independent third-party auditors, and to provide auditors with all documentation and information needed to perform an audit. This bill was ultimately vetoed by Governor Newsom.[7]

- SB 243 (Padilla; Ch. 677, Stats. 2025) requires companion chatbot operators to disclose that a chatbot is artificial if a reasonable person interacting with a companion chatbot would be misled to believe that the person is interacting with a human, and mandates that chatbot operators implement protocols to respond when a user expresses suicidal ideation or self-harm, including providing contact information for crisis or suicide hotlines. Version 95 of this bill would have required operators to submit their companion chatbot platforms to regular audits by an independent third party to ensure compliance and would have required an operator to make a high-level summary of the results of those audits publicly available at no cost to a person who accesses the operator's internet website.

- AB 2023 (Wicks, Bauer-Kahan) and SB 1119 (Padilla) are dual companion chatbot safety bills currently moving through the Legislature. These bills create an independent auditing framework wherein the Attorney General adopts regulations establishing professional standards for auditors, and operators of companion chatbots are required to annually submit to an independent audit assessing compliance with the bill's provisions. The Attorney General is permitted to disclose specified information from an audit report to a Government Agency or a public prosecutor, as appropriate for enforcement purposes.

Despite numerous attempts, no bill containing an independent auditing requirement for AI systems or models has yet been enacted in California. A letter of opposition submitted by the California Chamber of Commerce, TechNet, TechCA, and the Computer & Communications Industry Association for AB 1405 (Bauer-Kahan, 2025) offers a possible explanation for this phenomenon:

> What matters perhaps even more than auditors being "third party" evaluators or reviewers is that they are independent, standards-based and standards-bound professionals, with appropriate certification confirming their credentials demonstrating their expertise as well as their accountability, public trust and confidence via a framework that at minimum:
>
> - provides the necessary oversight that makes them subject to reputational and legal consequences for their misconduct, negligence, if not ethical violations;

[7] Office of the Governor, (Oct. 13, 2025), https://www.gov.ca.gov/wp-content/uploads/2025/10/AB-1064-Veto.pdf

- ensures their independence/impartiality/objectivity and affirms their duty of confidentiality to protect proprietary information, trade secrets, shareholder interests and other business information, unless legally compelled to disclose; and,

- establishes clear standards by which companies, shareholders, the public, and the courts can evaluate and rely upon their reports for oversight and enforcement.

It should be noted that today's AI auditing ecoysystem is virtually non-existent, lacking not only comprehensive standards in terms of standard audit procedures as well robust frameworks for governing the professional conduct of the AI auditors themselves, but also lacking sufficient resources for conducting invariably complex audits for AI systems.

These remarks suggest the primary obstacle to requiring independent oversight of AI systems and models is not a lack of interest in AI auditing, but rather the absence of a robust framework for determining who may conduct audits, what standards they must apply, and how auditor independence, competence, and accountability should be ensured.

Other states. Earlier this year, Illinois' legislature passed SB 315 (Rep. Didech, Sen. Edly-Allen) and sent it to Governor Pritzker for signing. SB 315 is a frontier-AI safety bill requiring developers of large foundation models with more than $500 million in annual gross revenue to create, follow, publish, and annually update a frontier AI framework addressing catastrophic-risk assessment, mitigations, cybersecurity, governance, and internal-use risks.[8] While largely mirroring California's Transparency in Frontier AI Act, SB 315 departs from California law in requiring large frontier developers to annually retain third parties to perform independent compliance audits. Third party auditors are required to "conduct audits consistent with generally accepted auditing standards and best practices" and are granted access to all materials reasonably necessary to comply with the bill. Reports produced by auditors must contain all of the following information:

- A description of whether the large frontier developer has substantially complied with the requirements of the bill.

- A description of material deviations from the requirements of the bill, an explanation of any deviation and its rationale, and any recommendations for how a developer can improve its policies and processes for ensuring compliance, if applicable.

- A detailed assessment of the large frontier developer's internal controls, including its designation and empowerment of senior personnel responsible for such implementation by the large frontier developer, its employees, and its contractors.

- A list of the personnel involved in the audit.

---

[8] Illinois SB 315, https://www.ilga.gov/Legislation/BillStatus/FullText?GAID=18&DocNum=315&DocTypeID=SB&LegId=157797&SessionID=114.

- The third party's procedures for managing conflicts of interest and any conflicts of interest of any personnel involved in the audit.

- The methodology of the audit and the nature of the information reviewed by the third party to conduct the audit.

- The signature of the lead auditor certifying the results of the audit.

The bill permits the Attorney General and Illinois Emergency Management Agency and Office of Homeland Security access to audit reports and requires large developers to conspicuously publish high-level summaries of audit findings on their websites. Governor Pritzker has signaled he intends to sign the bill, and large model developers OpenAI and Anthropic have come out publicly in support of the measure.[9]

Federal efforts. Earlier this month, Representatives Jay Obernolte (R-CA) and Lori Trahan (D-MA) released a "discussion draft" of a bill titled the Great American AI Act.[10] Among its various provisions, this bill would establish a Center for AI Standards and Innovation tasked with developing guidelines, best practices, and voluntary standards for covered entities to evaluate and improve security measures with respect to artificial intelligence systems, as well as with conducting evaluations and assessments with respect to AI systems developed in the United States or by foreign adversaries. The Center would also be tasked with licensing independent verification organizations (IVOs) – a concept shared with SB 813 – to perform auditing of frontier developers' compliance with the requirements of the Act. With respect to IVOs, the Center would additionally do the following:

- Establish qualifications, conflict-of-interest requirements, and standards of practice applicable to IVOs.

- Receiving and reviewing the results of IVO-conducted audits and assessments.

- Referring matters arising from audits and assessments to the Attorney General for enforcement action, or to other Federal agencies with relevant regulatory jurisdiction as appropriate.

- Publishing periodic reports summarizing aggregate audit and assessment findings.

- Suspending, revoking, or otherwise disciplining licensed IVOs for failure to comply with the requirements of the Act.

Under the Act, a large frontier developer is required to retain a licensed IVO semiannually to perform ongoing verification of the developer's compliance, as well as ongoing assessment of the adequacy of the developer's AI framework, governance policies and practices, risk-

---

[9] Jared Perlo, "Illinois Legislature passes historic AI bill that would require third-party safety audits," *NBC News,* (May 27, 2026), https://www.nbcnews.com/tech/tech-news/illinois-legislature-passes-historic-ai-bill-rcna347191.
[10] Jay Obernolte and Lori Trahan, "The Great American AI Act - Discussion Draft," https://obernolte.house.gov/sites/evo-subsites/obernolte.house.gov/files/evo-media-document/the-great-american-ai-act-discussion-draft-website-compressed-compressed.pdf

monitoring, and mitigation of detected risks. Developers are required to grant IVOs timely access to unredacted materials, personnel, and records.

Relationship to AB 1405. According to a press release from the author of this measure, SB 813 is meant to work in concert with AB 1405 (Bauer-Kahan, 2025):

> State Sen. Jerry McNerney, D-Pleasanton, and Assemblymember Rebecca Bauer-Kahan, D-Orinda, announced today that they are partnering to establish landmark safety standards for independent, third-party assessments of AI systems and models.
>
> . . .
>
> Sen. McNerney's SB 813 legislation would establish the California Artificial Intelligence Standards and Safety Commission made up of AI industry experts, academics, and public officials who would create voluntary AI safety standards across a range of industries and applications.
>
> Asm. Bauer-Kahan is planning to amend her AB 1405 to establish a robust registry for independent, third-party AI auditors to verify the safety of AI systems and models.

AB 1405 is currently pending on the Senate Appropriation Committee's suspense file.

The Business and Consumer Services Agency. In July of 2025, Governor Newsom announced the Business, Consumer Services, and Housing Agency would be divided into two new agencies: the Business and Consumer Services Agency (BCSA), and the Housing and Homelessness Agency.[11] According to the Governor's press release, BCSA will house the Department of Financial Protection and Innovation and Department of Consumer Affairs, among others, and will focus on "business regulation and consumer protection." In May of this year, Governor Newsom appointed Rohit Chora, former Director of the US Consumer Financial Protection Bureau and Commissioner of the Federal Trade Commission, as the Secretary of BCSA.[12] The agency will officially launch on July 1st of this year.

3) What this bill would do. This bill would create a seven-member California Artificial Intelligence Standards and Safety Commission in BCSA, and require the commission to (1) develop standards for the assessment of AI systems and models by AI auditors, (2) develop procedures for determining whether to suspend or terminate the registration of an AI auditor, and (3) develop criteria for determining whether an AI auditor qualifies as an independent verification organization.

The bill envisions two tiers of standards for any given AI system, model, or use case: a "minimum compliance standard" that establishes criteria for determining whether an AI system

---

[11] Governor's Press Office, "Governor Newsom restructures state government to combat homelessness, boost housing and affordability," (Jul. 11, 2025), https://www.gov.ca.gov/2025/07/11/governor-newsom-restructures-state-government-to-combat-homelessness-boost-housing-and-affordability/

[12] Governor's Press Office, "Governor Newsom appoints former federal regulator Rohit Chopra to head new Business and Consumer Services Agency amid Trump-era rollbacks," (May 12, 2026), https://www.gov.ca.gov/2026/05/12/governor-newsom-appoints-former-federal-regulator-rohit-chopra-to-head-new-business-and-consumer-services-agency-amid-trump-era-rollbacks/

or model meets minimum safety, efficacy, reliability, and security requirements necessary to comply with applicable state law, and an "advanced safety standard" that establishes criteria for determining whether an AI system or model substantially exceeds minimum compliance requirements, incorporates emerging best practices, and achieves demonstrable outcomes that promote the safety, efficacy, reliability, and security of the system or model. Advanced safety standards are intended to be applied by "independent verification organizations" (IVOs), AI auditors that specialize in assessing whether AI systems and models are designed according to industry best practices.

In the interest of not creating duplicative commissions, the definition of "commission" in this bill contains language that would confer authority over this bill's provisions onto the e-Safety Commission created by AB 1709 (Lowenthal, 2026), should that bill be enacted and the commission created.

Omidyar Network, writing in support, describes the need for this legislation:

> Artificial intelligence presents extraordinary opportunities, but the pace of technological development has far outstripped the pace of public oversight. As AI systems become increasingly powerful and integrated into daily life, policymakers must ensure that appropriate governance structures exist to protect consumers, promote innovation, and build public trust. SB 813 and AB 1405 represent the next logical step in California's leadership on AI governance by establishing the institutional infrastructure necessary to ensure that Californians can benefit from the promise of AI, while minimizing preventable harms associated with its misuse. The legislation creates a two-tiered, voluntary framework to evaluate AI systems and assess whether models meet advanced safety standards and reflect evolving best practices. These bills establish a benchmarking system in which the State of California sets desired public-interest outcomes—including safety, privacy, security, and accuracy—while independent, qualified, and accredited experts verify whether AI systems achieve those outcomes. This flexible, market-based and future-oriented approach reflects California's long-standing tradition of innovation while reinforcing our state's role as a global leader in consumer protection and responsible technology governance.

The Little Hoover Commission, writing in support, describes the importance of establishing a central entity to oversee AI policy:

> In its 2024 report, Artificial Intelligence and California State Government, the Commission found that artificial intelligence presents significant opportunities to improve government services and operational efficiency. But it noted that AI also carries substantial risks, including bias, misinformation, and governance failures if deployed without clear safeguards and coordinated oversight. To protect against these risks and ensure all of the state's AI efforts adhere to the core principles of transparency, accountability, and safety, the Commission recommended that California establish an AI Council to oversee the implementation of AI in state operations.

4) Committee amendments. Encode, striking a "support if amended" position, requests that the Commission proposed under the bill be amended to allow for shared appointment authority:

> The Legislature — through the Senate Committee on Rules and the Speaker of the Assembly — should be granted authority to appoint a meaningful share of Commission members (e.g.

> Cal. Gov. Section 11546.8 (g)). This would ensure no single branch of government exercises undue influence over a Commission tasked with shaping AI safety standards that will affect the public, industry, and government alike.

In response to this concern, the author has agreed to a committee amendment granting the Speaker of the Assembly and Senate Rules Committee one appointment each.

Adopting an "oppose unless amended" position, the California Chamber of Commerce describes issues with the bill related to transparency and capacity:

> AI is a general-purpose technology deployed across highly specialized sectors, including healthcare, financial services, critical infrastructure, transportation, cybersecurity, education, scientific research, and countless other domains. Developing meaningful safety, auditing, and risk-assessment standards across these diverse use cases requires substantial domain-specific expertise and sustained stakeholder participation. Established standards-development organizations are designed to convene thousands of technical experts, practitioners, researchers, and stakeholders across specialized committees over many years to develop consensus standards.
>
> By contrast, SB 813 would concentrate significant authority within a seven-member commission that would be expected to evaluate, adapt, oversee, and maintain standards affecting a broad range of industries and applications. No seven individuals can reasonably be expected to possess the collective expertise necessary to evaluate technical standards, auditing methodologies, and risk-assessment frameworks across all of these domains simultaneously.
>
> . . .
>
> We further note that while it appears that the commission will be setting standards for AI safety in California, SB 813 does not clearly specify the procedural framework governing adoption of the commission's standards [i.e., whether they need to go through the standard Administrative Procedure Act (APA) process], and whether those standards are intended to function as advisory guidance, de facto regulatory benchmarks, or prerequisites for compliance under other statutory or regulatory regimes.
>
> To the extent these standards are expected to influence regulated entities, additional clarity is warranted regarding how they intersect with APA requirements, including notice-and-comment procedures, evidentiary standards, and judicial review. Without such clarity, there is a risk that quasi-regulatory standards may be developed outside of established administrative safeguards that ensure transparency, accountability, and consistency in rulemaking.

To address these concerns, the author may wish to build additional staffing capacity into the commission either by incorporating it more fully into the existing BCSA structure, or by adopting a framework wherein the commission establishes working groups to develop individual standards and make recommendations to the full commission. The author may wish to tie the ultimate adoption of these recommendations to the APA process to ensure transparency and fairness.

To address further stakeholder feedback, the author has additionally agreed to committee amendments that do all of the following:

- Recognize that minimum compliance standards and advanced safety standards can be "created or otherwise authorized by statute."
- Simplify the definition of "covered audit."
- Recast the definition of "minimum compliance standard" in terms of "operational effectiveness" and "design and implementation of internal controls."
- Replace the requirement for a commission member to be "the president of the California Board of Accountancy or that person's designee" with a commission member who is "an individual authorized to practice public accountancy in this state, appointed by the Governor."
- Acknowledge the role of national and international auditing and assurance standard-setting organizations, professional accountancy bodies in setting standards for consideration by the commission.
- In developing criteria for an AI auditor to be designated as an IVO, require the commission to consider whether the author both employs competent personnel and is capable of managing conflicts of interest.
- Require the commission, to the extent practicable, to align with existing frameworks and minimize duplicative requirements.
- Recast a requirement for the commission to review work products "at least annually" with a requirement to review "regularly," in order that the commission not be in a constant state of review as the number of work products grows.
- Provide that in an action alleging that a defendant's development, modification, or use of an artificial intelligence system or model caused harm, the fact that an audit has been performed on the system or model in accordance with a standard published under this chapter is relevant to, but not conclusive of, the action.
- Make various minor technical and conforming changes based on stakeholder feedback throughout.

The full text of SB 813 as proposed to be amended follows:

SEC. 2. Chapter 14 (commencing with Section 8898) is added to Division 1 of Title 2 of the Government Code, to read:

CHAPTER 14. California Artificial Intelligence Standards and Safety Commission

8898. As used in this chapter:

(a) "Advanced safety standard" means a standard published by the agency***, or created or otherwise authorized by statute,*** that establishes criteria for determining whether an AI system or model substantially exceeds minimum compliance requirements, incorporates emerging best practices, and achieves demonstrable outcomes that promote the safety, efficacy, reliability, and security of the system or model.

(b) "Agency" means the Business and Consumer Services Agency.

(c) "Artificial intelligence" or "AI" 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.

(d) "AI auditor" means a person, partnership, academic institution, nonprofit, or corporation that conducts a covered audit on behalf of a third party.

(e) (1) Except as provided in paragraph (2), "commission" means the California Artificial Intelligence Standards and Safety Commission established pursuant to Section 8898.1.

(2) If AB 1709 of the 2025–26 Regular Session is enacted and establishes an e-Safety Advisory Commission, "commission" means the e-Safety Advisory Commission established pursuant to AB 1709 of the 2025–26 Regular Session.

(f) "Covered audit" means ~~either of the following:~~

~~(1) An audit of an AI system or model conducted to verify whether the AI system or model complies with applicable state law.~~

~~(2) An~~***an*** assessment of an AI system or model conducted in accordance with ***a minimum compliance standard or*** an advanced safety standard.

(g) "Independent verification organization" or "IVO" means an AI auditor that is designated by the agency as having demonstrated expertise in assessing whether AI systems or models substantially exceed minimum compliance requirements, incorporate emerging best practices, and achieve demonstrable outcomes that promote their safety, efficacy, reliability, and security.

(h) "Minimum compliance standard" means a standard published by the agency***,*** or created ***or otherwise authorized*** by statute***,*** that establishes criteria for determining whether ~~an AI system or model meets~~***both of the following meet*** minimum safety, efficacy, reliability, and security requirements necessary to ~~comply with~~ ***meet*** applicable ***requirements under California*** state law***:***~~.~~

(1) The operational effectiveness of an AI system or model.

(2) The design and implementation of internal controls established for an AI system or model.

8898.1. (a) On or before July 1, 2027, the agency shall establish the California Artificial Intelligence Standards and Safety Commission that consists of all of the following members:

(1) A representative, appointed by the Governor, from a business that develops or deploys AI systems or models and employs more than 100 workers.

(2) A representative, appointed by the Governor, from a business that develops or deploys AI systems or models and employs fewer than 100 workers.

(X) An individual authorized to practice public accountancy in this state, appointed by the Governor, who represents the accounting and auditing profession and has experience providing assurance services related to artificial intelligence and technology systems.

(3) An expert from academia, appointed by the Governor, with experience researching AI safety.

(4) An individual, appointed by the ~~Governor~~***Speaker of the Assembly***, representing the interests of civil society, including, but not limited to, nongovernmental organizations, public policy institutes, or consumer advocacy organizations.

(5) An individual, appointed by the ~~Governor~~***Senate Rules Committee***, representing the legal and economic interests of workers in relation to AI systems and models.

~~(6) The President of the California Board of Accountancy or that person's designee.~~

(7) The Attorney General or that person's designee.

(b) Members appointed pursuant to this subdivision shall serve four-year terms and may be reappointed.

(c) If AB 1709 of the 2025–26 Regular Session is enacted, and the e-Safety Advisory Commission is established pursuant to that measure, this section shall be inoperative.

8898.2. (a) The commission shall do all of the following:

(1) Develop standards for the assessment of AI systems and models by AI auditors. In developing these standards, the commission shall do all of the following:

(A) Evaluate existing standards, frameworks, guidelines, ***criteria,*** and best practices developed or published by government agencies, standards-setting organizations, ***including national and international auditing and assurance standard-setting organizations, professional accountancy bodies,*** AI auditors, private entities that develop or deploy AI systems or models, or other independent experts with relevant expertise.

(B) Establish the following two tiers of standards:

(i) Minimum compliance standards.

(ii) Advanced safety standards.

(C) Prioritize, at a minimum, developing standards for AI systems, models, and use cases that are subject to state law.

(D) Create a process for independent experts to submit advanced safety standards for evaluation by the commission pursuant to this chapter.

(2) Develop procedures for determining whether to suspend or terminate the registration of an AI auditor. In developing these procedures, the commission shall consider all of the following:

(A) Failures to adhere to appropriate standards, including standards developed pursuant to this chapter.

(B) Material misrepresentations in audit reports or required disclosures.

(C) Conflicts of interest that impair independence.

(D) Failure to maintain adequate documentation.

(E) Conduct that reasonably calls into question the integrity, objectivity, or competence of the AI auditor.

(3) Develop criteria for determining whether an AI auditor qualifies as an independent verification organization. In developing these criteria, the commission shall consider whether the IVO is capable of all of the following with respect to an AI system or model in operation:

(A) Assessing the safety, efficacy, reliability, security, and robustness of the AI system or model.

(B) Identifying reasonably foreseeable risks arising from the development, deployment, operation, use, or misuse of the AI system or model.

(C) Verifying that the AI system or model has achieved and will maintain a level of risk that is appropriate in light of the AI system's or model's intended purpose, foreseeable uses and misuses, affected populations, and the severity and likelihood of potential harms.

(D) Determining whether the AI system or model should be subject to ongoing monitoring, periodic reassessment, or other postdeployment review.

(E) Employing or otherwise engaging personnel with sufficient technical expertise to conduct covered audits.

(F) Identifying and managing potential conflicts of interest that may undermine the integrity, quality, or independence of covered audits.

(x) To the extent practicable, do both of the following:

(A) Align any standards, procedures, and criteria developed pursuant to this section with existing professional and regulatory audit and assurance standards, frameworks, and criteria to promote alignment, consistency, comparability, and reliability across regulatory and industry frameworks.

(B) Structure the requirements of this section to minimize duplicative compliance obligations, including by allowing reports, assessments, audits, or assurance engagements prepared to satisfy substantially similar requirements to be used for purposes of this section, provided such reports, assessments, audits, or engagements satisfy the requirements of this section.

(4) Publish any standards, procedures, and criteria developed pursuant to this chapter in a publicly accessible format on the agency's internet website. For each standard published pursuant to this paragraph, the commission shall identify, to the extent applicable, all of the following:

(A) Any state law to which the standard is intended to apply.

(B) Whether the standard is a minimum compliance standard or an advanced safety standard.

(C) The category, use case, or type of AI system or model to which the standard is intended to apply.

(D) The date upon which the standard was adopted and the date of the most recent revision to the standard.

(E) Any additional information necessary to facilitate the interpretation and application of the standard.

(5) Publish a statement on the agency's internet website prominently disclosing that the publication of a standard pursuant to this subdivision does not constitute recommendation or endorsement by the state of any AI system or model.

(6) ***Regularly*** ~~Review~~ ***review*** standards, procedures, and criteria published pursuant to this chapter ~~at least annually~~ and revise them as appropriate to reflect changes in applicable state law, technological developments, widely recognized industry standards, and emerging best practices.

(b) In developing ~~advanced safety~~ standards pursuant to this chapter, the commission shall incorporate ***objective and clearly defined*** criteria for assessing all of the following with respect to an AI system or model in operation:

(1) The safety, efficacy, reliability, security, and robustness of the AI system or model.

(2) Reasonably foreseeable risks arising from the development, deployment, operation, use, or misuse of the AI system or model.

(3) The adequacy of any measures adopted to prevent, mitigate, and respond to those risks.

(4) Whether the AI system or model has achieved and maintained a level of risk that is appropriate in light of the AI system's intended purpose, foreseeable uses and misuses, affected populations, and the severity and likelihood of potential harms.

(5) Whether the AI system or model should be subject to ongoing monitoring, periodic reassessment, or other postdeployment review.

(6) Whether a developer, deployer, or other responsible entity has implemented and adhered to appropriate internal safety standards and protocols.

(7) Any other criteria specified in applicable state law.

8898.3. (a) In carrying out its duties under this chapter, the commission ~~may~~ ***shall*** consult ~~and coordinate~~ with any of the following ***as appropriate***:

(1) Certified public accountants and public accountants in good standing as certified by the California Board of Accountancy.

(2) AI auditors that are not certified public accountants or public accountants.

(3) Independent verification organizations.

(4) Academic institutions.

(5) Private entities that develop or deploy AI systems or models.

(6) Consumer protection, labor, and civil society organizations.

(7) Relevant federal, state, and local agencies.

(8) National and international standards-setting organizations.

(9) Any other relevant stakeholders, as determined by the commission.

(b) The commission shall convene working groups to solicit stakeholder input in the development and revision of standards, procedures, and criteria pursuant to this chapter.

8898.4. The commission may establish voluntary pilot programs to evaluate the effectiveness of standards developed pursuant to this chapter.

8898.5. This chapter does not do either of the following:

(a) Establish liability solely for failure to comply with a standard developed or published pursuant to this chapter.

(b) Constitute recommendation or endorsement by the state of any AI system or model.

8898.X. In an action alleging that a defendant's development, modification, or use of an artificial intelligence system or model caused harm, the fact that an audit has been performed on the system or model in accordance with a standard published under this chapter is relevant to, but not conclusive of, the action.

8898.6. This chapter shall become operative only if Assembly Bill 1405 of the 2025–26 Regular Session is enacted.

ARGUMENTS IN SUPPORT: The Alliance for Secure AI writes in support:

> Right now, there are no consistent requirements for safety testing and evaluation in advanced AI, and no reliable standards for who is qualified to conduct that testing. As these systems grow more capable, both gaps become harder to defend. We need a baseline for who can credibly evaluate AI systems and what those evaluations should measure. Together, SB 813 and AB 1405 build that infrastructure, and we encourage the committee to advance them.
>
> . . .
>
> SB 813 builds on that foundation with the Independent Verification Organization model, and it gets something right that many AI proposals miss. Rather than writing technical requirements into statute, where they risk becoming obsolete before implementation, SB 813 creates a standing body, the California Artificial Intelligence Standards and Safety Commission, to develop, update, and publish safety standards through independent experts. The Commission sets the standards and the criteria for verification; independent auditors and verification organizations apply them. It also separates a baseline of compliance from an advanced safety tier, giving developers who exceed the minimum a credible way to show it. That division of labor reflects how serious safety regimes work, and it gives California a framework that can evolve alongside the technology.

ARGUMENTS IN OPPOSITION: Technet writes in opposition:

> AI development and deployment are inherently global, and a state-specific definition of "substantially exceeds compliance" adds to a growing patchwork rather than reducing it. Other states are pursuing similar but distinct frameworks — including an IVO proposal in Ohio (HB 628) and a pilot program in Connecticut (SB 5) — each developed through different processes and reflecting different assumptions. AI developers and deployers operating across state lines would face multiple, potentially conflicting sets of "best practices" that may be technically incompatible with one another, complicating compliance without a clear corresponding improvement in safety outcomes.
>
> We also note that California is concurrently negotiating audit and assessment requirements for frontier AI models through SB 1119 and AB 2023. Standing up a separate, two-tier auditing framework under SB 813 while those negotiations are ongoing risks creating two parallel and potentially inconsistent compliance regimes for the same underlying technology. California has historically played a leadership role by aligning with broader standards-setting efforts rather than creating siloed frameworks that diverge from national and international approaches, and we would encourage that same approach here.

REGISTERED SUPPORT / OPPOSITION:

Support

Fathom (sponsor)
Children Now
Eticas
Little Hoover Commission
Meridian Governance INC.
Mothers Against Media Addiction

Oakland Privacy
Omidyar Network
The Alliance for Secure Ai
Transparency Coalition.ai
4 Academics

Support If Amended

Encode Ai Corporation
Secure Ai Project

Oppose

Technet
Valley Industry and Commerce Association (VICA)

Oppose Unless Amended

California Chamber of Commerce

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