SENATE JUDICIARY COMMITTEE
Senator Thomas Umberg, Chair
2025-2026 Regular Session

SB 53 (Wiener)
Version: March 27, 2025
Hearing Date: April 8, 2025
Fiscal: Yes
Urgency: No
CK

SUBJECT

CalCompute: foundation models: whistleblowers

DIGEST

This bill establishes a consortium tasked with developing a framework for a public cloud computing cluster that advances the ethical development and deployment of AI for the public good. This bill creates protections for whistleblowers working with specified AI models when reporting on "critical risks" and requires developers to provide processes for anonymous reporting of activities posing such risks.

EXECUTIVE SUMMARY

Large AI models have opened beneficial possibilities and breakthroughs in a variety of sectors, including healthcare, creative industries, education, business operations, and customer service. However, these models also pose significant risks to society, whether it is creation of chemical, biological, radiological, or nuclear (CBRN) weapons, utilization to carry out cyberattacks at scale, or evasion of the oversight and control of their developers or deployers. Given how complex and opaque these models can be and the lack of comprehensive regulatory oversight of their development, the need for stronger protections for experts working on these models who come forward to warn of such risks is clear.

This bill establishes baseline whistleblower protections for employees that disclose information to proper authorities or management where there is reasonable cause to believe that the information discloses activities on the part of a developer that poses a "critical risk" or that discloses misstatements about the management of critical risks. Developers are required to notify employees of these protections and to create internal reporting systems for anonymous disclosures of this sort. The bill also establishes a consortium tasked with developing a framework for a public cloud computing cluster, to be called "CalCompute," that will advance ethical AI development for public benefit to counter the imbalance in access to AI infrastructure.

The bill is sponsored by Encode, Secure AI Project, and Economic Security California Action. It is supported by a number of advocacy organizations, including Oakland Privacy. No timely opposition was received by the Committee. The bill passed out of the Senate Governmental Organization Committee on a vote of 13 to 0.

PROPOSED CHANGES TO THE LAW

Existing law:

1) Prohibits employers and any person acting on behalf of the employer from making, adopting, or enforcing a rule, regulation, or policy preventing an employee from disclosing information to certain entities or from providing information to, or testifying before, any public body conducting an investigation, hearing, or inquiry if the employee has reasonable cause to believe that the information discloses a violation of a law, as specified. Employers and their agents are also prohibited from retaliating against an employee for such conduct. (Lab. Code § 1102.5.)

2) Requires the office of the Attorney General to maintain a whistleblower hotline to receive calls from persons who have information regarding possible violations of state or federal statutes, rules, or regulations, or violations of fiduciary responsibility by a corporation or limited liability company to its shareholders, investors, or employees. The Attorney General is required to refer calls received on the whistleblower hotline to the appropriate government authority for review and possible investigation. During the initial review of such a call, the Attorney General or appropriate government agency shall hold in confidence information disclosed through the whistleblower hotline, including the identity of the caller disclosing the information and the employer identified by the caller. (Lab. Code § 1102.7.)

3) Establishes the California Department of Technology (CDT) within the Government Operations Agency (GovOps), which is tasked with conducting, in coordination with other interagency bodies as it deems appropriate, a comprehensive inventory of all high-risk automated decision systems that have been proposed for use, development, or procurement by, or are being used, developed, or procured by, any state agency. (Gov. Code §§ 11545(a), 11546.45.5.)

4) Provides that the perfecting of an appeal stays proceedings in the trial court upon the judgment or order appealed from or upon the matters embraced therein or affected thereby, including enforcement of the judgment or order, but the trial court may proceed upon any other matter embraced in the action and not affected by the judgment or order. When there is a stay of proceedings other than the enforcement of the judgment, the trial court shall have jurisdiction of proceedings related to the enforcement of the judgment as well as any other

matter embraced in the action and not affected by the judgment or order appealed from. (Code Civ. Proc. § 916.)

This bill:

1) Establishes within GovOps a consortium to develop a framework for the creation of a public cloud computing cluster to be known as "CalCompute" that advances the development and deployment of AI that is safe, ethical, equitable, and sustainable by doing, at a minimum, both of the following:
    a) Fostering research and innovation that benefits the public.
    b) Enabling equitable innovation by expanding access to computational resources.

2) Provides that the consortium shall operate in accordance with all relevant labor and workforce laws and standards and make reasonable efforts to ensure that CalCompute is established within the University of California (UC) to the extent possible.

3) Requires CalCompute to include, but not be limited to, all of the following:
    a) A fully owned and hosted cloud platform.
    b) Necessary human expertise to operate and maintain the platform.
    c) Necessary human expertise to support, train, and facilitate the use of CalCompute.

4) Requires GovOps, on or before January 1, 2027, to submit a report from the consortium to the Legislature with the framework for CalCompute. The report shall include all of the following elements:
    a) A landscape analysis of California's current public, private, and nonprofit cloud computing platform infrastructure.
    b) An analysis of the cost to the state to build and maintain CalCompute and recommendations for potential funding sources.
    c) Recommendations for the governance structure and ongoing operation of CalCompute.
    d) Recommendations for the parameters for use of CalCompute, including, but not limited to, a process for determining which users and projects will be supported by CalCompute.
    e) An analysis of the state's technology workforce and recommendations for equitable pathways to strengthen the workforce, including the role of CalCompute.
    f) A detailed description of any proposed partnerships, contracts, or licensing agreements with nongovernmental entities, including, but not limited to, technology-based companies, that demonstrates compliance with the requirements of subdivisions (c) and (d).

   g) Recommendations regarding how the creation and ongoing management of CalCompute can prioritize the use of the current public sector workforce.

5) Provides that the consortium shall, consistent with state constitutional law, consist of 14 members selected from specified fields and industries. Eight members are to be selected by GovOps, and the President pro Tempore of the Senate and the Speaker of the Assembly shall each select three members. The members of the consortium shall serve without compensation, but shall be reimbursed for all necessary expenses actually incurred in the performance of their duties. The consortium shall be dissolved upon submission of the report.

6) Provides that if CalCompute is established within the UC, the UC may receive private donations for the purposes of implementing CalCompute.

7) Provides that the above provisions shall become operative only upon an appropriation in a budget act, or other measure, for the purposes of this section.

8) Prohibits a developer from making, adopting, or enforcing a rule, regulation, or policy that prevents an employee from disclosing, or retaliates against an employee for disclosing, information to the Attorney General, federal authorities, or another employee who has authority to investigate, discover, or correct the reported issue, if the employee has reasonable cause to believe that the information discloses either of the following:
   a) The developer's activities pose a critical risk.
   b) The developer has made false or misleading statements about its management of critical risk.

9) Defines "critical risk" as a foreseeable and material risk that a developer's development, storage, or deployment of a foundation model will result in the death of, or serious injury to, more than 100 people or more than $1 billion in damage to rights in money or property, through any of the following:
   a) The creation and release of a CBRN weapon.
   b) A cyberattack.
   c) A foundation model engaging in conduct, with limited human intervention, that would, if committed by a human, constitute a violation of the Penal Code that requires intent, recklessness, or gross negligence or the solicitation or aiding and abetting of that violation.
   d) A foundation model evading the control of its developer or user.

10) Authorizes an employee to use the hotline described in Section 1102.7 of the Labor Code to make the above reports.

11) Requires a developer to provide a clear notice to all employees of their rights and responsibilities hereunder, as specified.

12) Requires a developer to provide a reasonable internal process through which an employee may anonymously disclose information to the developer if the employee believes in good faith that the information indicates that the developer's activities present a critical risk, including a monthly update to the person who made the disclosure regarding the status of the developer's investigation of the disclosure and the actions taken by the developer in response to the disclosure.

13) Requires the disclosures and responses associated with the above process to be shared with officers and directors of the developer at least once each quarter, except as to any officer or director accused of wrongdoing.

14) Authorizes civil actions to be brought for violations. The court is authorized to award reasonable attorney's fees and injunctive relief to a plaintiff who brings a successful action for a violation of this section. In a civil action brought pursuant hereto, once it has been demonstrated by a preponderance of the evidence that an activity proscribed was a contributing factor in the alleged prohibited action against the employee, the developer shall have the burden of proof to demonstrate by clear and convincing evidence that the alleged action would have occurred for legitimate, independent reasons even if the employee had not engaged in activities protected by this section.

15) Provides that, in a civil action or administrative proceeding brought pursuant hereto, an employee may petition the superior court in any county wherein the violation in question is alleged to have occurred, or wherein the person resides or transacts business, for appropriate temporary or preliminary injunctive relief. Upon the filing of the petition for injunctive relief, the petitioner shall cause notice thereof to be served upon the person, and thereupon the court shall have jurisdiction to grant temporary injunctive relief as the court deems just and proper.

16) Provides that, in addition to any harm resulting directly from a violation of this section, the court shall consider the chilling effect on other employees asserting their rights under this section in determining whether temporary injunctive relief is just and proper. Appropriate injunctive relief shall be issued on a showing that reasonable cause exists to believe a violation has occurred.

17) Requires an order authorizing temporary injunctive relief to remain in effect until an administrative or judicial determination or citation has been issued, or until the completion of a review, as provided, or at a certain time set by the court. Thereafter, a preliminary or permanent injunction may be issued if it is shown to

be just and proper. Any temporary injunctive relief shall not prohibit a developer from disciplining or terminating an employee for conduct that is unrelated to the claim of the retaliation. Notwithstanding Section 916 of the Code of Civil Procedure, injunctive relief granted pursuant hereto shall not be stayed pending appeal.

COMMENTS

1. Whistleblower protections

The powerful benefits and existential risks of AI technology, and particularly large generative AI models, is well documented. Despite a flurry of introduced legislation across the country, there is still not a comprehensive regulatory framework to ensure the safe and equitable development of this technology. This bill seeks to implement baseline protections for potential whistleblowers to ensure those closest to the development of this technology feel comfortable bringing concerns to the fore.

Whistleblower protections within companies developing AI foundation models are crucial given the extraordinary scale and impact of the technology. These protections serve as an essential safety valve in an industry where the potential consequences of unethical or dangerous development practices can affect millions, if not billions, of people.

These models increasingly power critical infrastructure across healthcare, finance, employment, education, and government services. When employees with direct knowledge of risks, harms, or unethical practices cannot safely speak up, dangerous systems may be deployed without proper safeguards or public awareness. Whistleblowers often represent the last line of defense when corporate incentives prioritize growth, profit, or competitive advantage over public welfare.

The technical complexity of foundation models creates significant information asymmetry between companies and outside regulators or the public. Most governmental bodies lack the specialized expertise to effectively evaluate these systems without insider information. This knowledge gap makes whistleblowers uniquely positioned to identify issues that might otherwise remain hidden from external reviewers. Their technical understanding allows them to recognize significant risks that non-specialists or outsiders might miss.

Given the regulatory vacuum in this industry, internal voices raising concerns about potential harms serve a crucial function in preventing the deployment of unsafe systems. Without protected channels for these voices, companies may face fewer checks on irresponsible development practices until after harm has occurred.

This bill prohibits a "developer" from making, adopting, or enforcing a rule, regulation, or policy that prevents an employee from disclosing, or retaliates against an employee for disclosing, information to the Attorney General, federal authorities, or another employee who has authority to investigate, discover, or correct the reported issue, if the employee has reasonable cause to believe that the information discloses that the developer's activities pose a critical risk or that the developer has made false or misleading statements about its management of critical risk.

"Developer" is narrowly defined to include only those persons who have trained a foundation model with a quantity of computational power that costs at least $100 million. "Critical risk" is defined to mean a foreseeable and material risk that a developer's development, storage, or deployment of a foundation model will result in the death of, or serious injury to, more than 100 people or more than $1 billion in damage to rights in money or property, through any of the following:

- The creation and release of a CBRN weapon.
- A cyberattack.
- A foundation model engaging in conduct, with limited human intervention, that would, if committed by a human, constitute a violation of the Penal Code that requires intent, recklessness, or gross negligence or the solicitation or aiding and abetting of that violation.
- A foundation model evading the control of its developer or user.

Developers are required to provide notice to all employees of their rights and responsibilities hereunder, including posting and displaying a notice to all employees of their rights, ensuring that any new employee receives equivalent notice, and ensuring that any employee who works remotely periodically receives an equivalent notice; or annually providing written notice to each employee of the employee's rights hereunder and ensuring that the notice is received and acknowledged by all of those employees.

A developer must also establish a reasonable internal process for employees to anonymously disclose information to the developer if the employee believes in good faith that the information indicates that the developer's activities present a critical risk. The process must include a monthly update to the person who made the disclosure regarding the status of the developer's investigation of the disclosure and the actions taken by the developer in response to the disclosure.

Violations may be remedied through civil or administrative actions and courts are authorized to award a prevailing plaintiff injunctive relief and reasonable attorney's fees. In a civil action, the bill alters the relevant burdens. Once it has been demonstrated by a preponderance of the evidence that a relevant protected activity was a contributing factor in the alleged prohibited action against the employee, the developer shall have the burden of proof to demonstrate by clear and convincing evidence that the alleged

action would have occurred for legitimate, independent reasons even if the employee had not engaged in the protected activities. Appropriate injunctive relief shall be issued on a showing that reasonable cause exists to believe a violation has occurred.

2. CalCompute

As industry races toward developing larger, more powerful AI models and seeks to commodify the seemingly infinite applications of AI, concerns are growing about the diminishing role that researchers, academic institutions, and more public-focused entities are playing in the development of AI. As reported by the Washington Post:

> As such tech behemoths as Meta, Google and Microsoft funnel billions of dollars into AI, a massive resources gap is building with even the country's richest universities. Meta aims to procure 350,000 of the specialized computer chips — called GPUs — that are essential to run the gargantuan calculations needed for AI models. In contrast, Stanford's Natural Language Processing Group has 68 GPUs for all of its work.
>
> To obtain the expensive computing power and data required to research AI systems, scholars frequently partner with tech employees. Meanwhile, tech firms' eye-popping salaries are draining academia of star talent.
>
> Big tech companies now dominate breakthroughs in the field. In 2022, the tech industry created 32 significant machine learning models, while academics produced three, a significant reversal from 2014, when the majority of AI breakthroughs originated in universities, according to a Stanford report.
>
> Researchers say this lopsided power dynamic is shaping the field in subtle ways, pushing AI scholars to tailor their research for commercial use. Last month, Meta CEO Mark Zuckerberg announced that the company's independent AI research lab would move closer to its product team, ensuring "some level of alignment" between the groups, he said.
>
> "The public sector is now significantly lagging in resources and talent compared to that of industry," said [Fei-Fei] Li, a former Google employee and the co-director of the Stanford Institute for Human-Centered AI. "This will have profound consequences because industry is focused on developing technology that is profit-driven, whereas public-sector AI goals are focused on creating public goods."
> …
> As Silicon Valley races to build chatbots and image generators, it is drawing would-be computer science professors with high salaries and the chance to work on interesting AI problems. Nearly 70 percent of people

> with PhDs in AI end up in private industry compared with 21 percent of graduates two decades ago, according to a 2023 report.[1]

The bill seeks to address this by establishing a consortium with GovOps tasked with developing a framework for the creation of a public cloud computing cluster to be known as "CalCompute." The explicit direction is for CalCompute to advance the development and deployment of AI that is safe, ethical, equitable, and sustainable by, among other things, fostering research and innovation that benefits the public and enabling equitable innovation by expanding access to computational resources.

The bill requires that CalCompute include a fully owned and hosted cloud platform and the necessary human expertise to operate and maintain the platform and to support, train, and facilitate the use of it. A report shall be submitted to the Legislature by January 1, 2027 at which point the consortium will be dissolved.

3. Stakeholder positions

According to the author:

> The greatest innovations happen when our brightest minds have the resources they need and the freedom to speak their minds. California's leadership on AI is more critical than ever as the new federal Administration proceeds with shredding the guardrails meant to keep Americans safe from the known and foreseeable risks that advanced AI systems present.
>
> Senate Bill 53 recognizes that we need a multipronged approach to encourage responsible AI innovation. It does so by:
>
> 1. Ensuring employees of frontier model labs are not retaliated against for reporting their belief that a developer's activities pose a critical risk, as defined, or has made false or misleading statements about its management of critical risk; and
>
> 2. Establishing a process to create a public cloud-computing cluster (CalCompute) that will conduct research into the safe and secure deployment of large-scale artificial intelligence (AI) models.

[1] Naomi Nix, Cat Zakrzewski & Gerrit De Vynck, *Silicon Valley is pricing academics out of AI research* (March 10, 2024) The Washington Post, https://www.washingtonpost.com/technology/2024/03/10/big-tech-companies-ai-research/ [as of Apr. 2, 2025].

> In doing this, SB 53 allows California to continue to lead in this space and to demonstrate that safety does not stifle success.

A coalition of organizations, including the three co-sponsors, Secure AI Project, Encode, and Economic Security California Action, write:

> **AI researchers and engineers will likely be the first to know if a company's product poses a significant risk to public safety, but company contracts may prevent them from reporting concerns.** In general, companies offer or require broad non-disclosure and non-disparagement agreements that are quite broad, often limiting nearly all disclosures that are not explicitly protected by law. These agreements in practice can prevent employees from sharing a good faith public safety concern with a government authority, which could chill employees' disclosure rights when that is their best or only course of action.
>
> **An anonymous internal process would encourage whistleblowers to report safety issues when necessary and without the government needing to get involved.** But employees at AI companies in today's hypercompetitive markets may be concerned about reporting their concerns to their supervisors for fear of dismissal or retaliation. By requiring companies to have an internal process where employees can report dangers, SB 53 would improve information flow and in turn enhance companies' awareness of serious dangers. Many companies already have internal processes, such as OpenAI, Anthropic, xAI, and Microsoft, so this requirement would codify what is an emerging industry norm.
>
> **SB 53 would also carefully expand existing whistleblower protections, which is essential for the AI industry.** If internal reporting fails to adequately address an identified risk, employees need an option to turn to government authorities. In more established industries, safety hazards are likely to violate an existing law or regulation, and disclosure of those hazards would be protected under existing law. However, as a nearly unregulated industry, dangerous behavior at AI companies might not today be illegal and disclosure is not protected in the same manner. SB 53 addresses this imbalance by expanding essentially the same protections from retaliation already found for reporting violations of the law to material risks of serious catastrophic harm, as well as false statements about company safety practices.

SUPPORT

Economic Security California Action (sponsor)
Encode (sponsor)
Secure AI Project (sponsor)
AI Policy Tracker
Apart Research
Beri
Center for AI and Digital Policy
Center for AI Policy
Center for Human-Compatible AI
EarningsStream LLC
Indivisible CA: StateStrong
Nonlinear
Oakland Privacy
Porn Free Colorado
Redwood Research
Safe AI Future

OPPOSITION

None received

RELATED LEGISLATION

Pending Legislation: AB 1405 (Bauer-Kahan, 2025) establishes an enrollment process for AI auditors within GovOps. It creates a publicly accessible repository of AI auditors and requires that they adhere to minimum standards of transparency, confidentiality, and ethical conduct. It also provides for whistleblower protections in certain cases for employees employed by enrolled auditors. AB 1405 is currently in the Assembly Appropriations Committee.

Prior Legislation: SB 1047 (Wiener, 2024) would have required developers of powerful AI models and those providing the computing power to train such models to put appropriate safeguards and policies into place to prevent critical harms. It would have established a state entity to oversee the development of these models and called for the creation of a consortium to develop a framework for a public cloud computing cluster, as provided for by this bill.

Governor Newsom vetoed SB 1047, stating in part: "By focusing only on the most expensive and large-scale models, SB 1047 establishes a regulatory framework that could give the public a false sense of security about controlling this fast-moving technology. Smaller, specialized models may emerge as equally or even more

dangerous than the models targeted by SB 1047 - at the potential expense of curtailing the very innovation that fuels advancement in favor of the public good.

Adaptability is critical as we race to regulate a technology still in its infancy. This will require a delicate balance. While well-intentioned, SB 1047 does not take into account whether an AI system is deployed in high-risk environments, involves critical decision-making or the use of sensitive data. Instead, the bill applies stringent standards to even the most basic functions - so long as a large system deploys it. I do not believe this is the best approach to protecting the public from real threats posed by the technology."

PRIOR VOTES:

Senate Governmental Organization Committee (Ayes 13, Noes 0)

****