



May 8, 2024

TO: Members, Senate Appropriations Committee

**SUBJECT: SB 1047 (WIENER) SAFE AND SECURE INNOVATION FOR FRONTIER ARTIFICIAL INTELLIGENCE MODELS ACT
OPPOSE – AS AMENDED APRIL 30, 2024
SCHEDULED FOR HEARING – MAY 16, 2024**

The undersigned organizations must respectfully **OPPOSE SB 1047 (Wiener)** as amended April 30, 2024, which would enact the Safe and Secure Innovation for Frontier Artificial Intelligence Models Act to require frontier AI developers to make a positive safety determination before initiating training of a covered model, among other things. While we share your goal of ensuring the safe and responsible development of AI, we believe that it is an issue that is appropriately being addressed at the federal level and are concerned that **SB 1047** will add more confusion to the already-fragmenting AI regulatory landscape in the U.S.

In addition to creating inconsistencies with federal regulations, the bill demands compliance with various vague and impractical, if not technically infeasible, requirements for which developers will be subject to harsh penalties, including potential criminal liability. We are concerned that the bill regulates AI technology as opposed to its high-risk applications, creates significant regulatory uncertainty and therefore high compliance costs, and poses significant liability risks to developers for failing to foresee and block any harmful use of their models by others – all of which inevitably discourages economic and technological innovation. This, unfortunately, does not better protect Californians. Instead, by hamstringing businesses from developing the very AI technologies that could protect them from dangerous models developed in territories beyond California's control, it risks only making them *more* vulnerable.

SB 1047's shift from "positive safety determinations" to "limited duty exemptions" is largely a shift in terminology, not in the underlying public policy

Whereas prior iterations of **SB 1047** *required* a developer to determine if they can make a "*positive safety determination*" with respect to a frontier model prior to initiating training of the model, recent amendments to **SB 1047** instead *authorize* a developer to determine whether the model qualifies for a "*limited duty exemption*" prior to initiating training of that model. Yet, the shift from mandatory "positive safety determinations" to permissive "limited duty exemptions" is a distinction without a difference, particularly as the definition of a "limited duty exemption" in the amended version of the bill mirrors the definition of a "positive safety determination" in earlier iterations.

Furthermore, until a determination is made that the model is the qualifies for a "limited duty exemption," to exclude the possibility that a model has a "hazardous capability" (or come close to having one), the developer must comply with the exact same requirements as a developer who could not make a positive safety determination. These include, for example, implementing the capability to promptly enact a full shutdown of the covered model until the developer can make the applicable determination (positive safety determination or limited duty exemption). Thus, while the terminology has shifted, neither the elements in

making those determinations, nor the consequences of failing to make such a determination, are any different.

SB 1047 mandates compliance with novel requirements based on standards that are often overbroad, vague, and impractical, if not simply infeasible

At its core, **SB 1047** seeks to regulate frontier AI developers from innovating AI models that will result in any kind of foreseeable harm—even harms that would not manifest from the model itself. In doing so, the bill requires developers to comply with incredibly vague, broad, impractical, if not impossible, standards when developing “covered models”.

For example, **SB 1047** applies to AI models that either: (1) meet a size threshold (a computing power greater than 10^{26} integer or floating-point operations in 2024), or (2) that perform similarly. What the latter category of covered models looks like, however, is not entirely clear. (See Proposed Section 22602(f).) The bill merely states that they are models “trained using a quantity of computing power *sufficiently large that it could be reasonably expected to have similar or greater performance* as an AI model trained using a quantity of computing power greater than 10^{26} integer or floating-point operations in 2024 as assessed using benchmarks commonly used to quantify the general performance of state-of-the-art foundation models”. There is little to no certainty as to what this translates to in practice and, in any case, such thresholds will become obsolete within a year, requiring the law to change yet again. Moreover, by equating model size to risk, the definition of “covered models” is simultaneously overly broad and too narrow as smaller and/or less performant models can present much greater risks than large/higher performant ones. As a result, the bill both fails to adequately address the very real risks posed by small but malicious models and imposes significant costs on innovating performant but responsible ones.

By way of another example, **SB 1047’s** definition of “hazardous capability” is so broad that it in fact captures not only covered models that have the capability to be used to enable certain harms (e.g. the creation or use of a chemical, biological, radiological, or nuclear weapon, or other threats of “comparable severity”) in a way that would be significantly more difficult to cause without access to the covered models, but also those that have such capabilities, “even if the hazardous capability *would not manifest but for* fine tuning and posttraining modification performed by third-party experts intending to demonstrate those abilities” – meaning, if third parties essentially jailbreak the model. (Proposed Section 22602(n)(2).) The overbreadth of the definition aside, what is considered sufficiently *close* to possessing a hazardous capability to prevent a model from qualifying for a limited duty exemption, however, is even more unclear. Also unclear is what would be considered a “reasonable margin for safety”, or an unreasonable one.

And finally, **SB 1047** also requires that a developer “incorporate all applicable covered guidance” before making determining if a covered model qualifies for a limited duty exemption. However, industry and others are still trying to ascertain how to define what constitutes a highly-capable, foundational model and it is therefore unclear what will qualify as “industry best practices” for the purpose of incorporating all applicable covered guidance. Such regulatory uncertainty will inevitably discourage economic and technological innovation. It would make far more sense to let the NIST (the National Institute of Standards and Technology) complete its work first, after which safety and security protocols tied to those safety standards could be considered.

SB 1047 focuses exclusively on developer liability, deters open source development, imposes questionable requirements on operators of “computing clusters” and imputes harsh penalties

There are a host of other issues and unintended consequences that warrant further consideration:

- **SB 1047 fails to account for the AI value chain, impeding open source.** The bill almost exclusively focuses on developer liability, failing to account for the AI value chain. Under **SB 1047**, developers must build full shutdown capabilities into their models and may be held liable for downstream uses over which they have no control, impeding their ability to open-source their models. Ultimately, liability should rest with the user who intended to do harm, as opposed to automatically defaulting to the developer who could not foresee, let alone block, any and all conceivable uses of a model that might do harm.

- **SB 1047 sets unreasonable safety incident reporting requirements that are not only vague but deter open-source development.** Developers are required to report each AI safety incident upon learning of it, or learning facts that would lead to the reasonable belief that a safety incident occurred. However, what is considered an “AI safety incident” is vague. Among other things, it includes a covered model “autonomously engaging in a sustained sequence of unsafe behavior other than at the request of a user” but fails to define what is considered “unsafe”, leaving developers to guess if they must report an incident. At the same time, “AI safety incident” covers a range of circumstances that are incompatible with open source because it would require monitoring of all downstream uses and applications.
- **SB 1047 imposes intrusive, if not unreasonable, requirements on operators of “computing clusters”.** Under the bill, there are a host of requirements that apply to any company that “operates a computing cluster” – presumably, data centers or cloud computing companies that provide cloud compute for frontier model training. As drafted, however, it is unclear as to what the bill means by “operate”, given that several entities could technically be seen operating a computer cluster: the owner of the cluster, the owner of the software operating the cluster, or the owner of the instance operating the cluster.

Moreover, the bill not only forces operators of computing clusters to collect personally identifiable data from their prospective customers, but it expects them to predict if a prospective customer “intends to utilize the computing cluster to deploy a covered model,” and requires that they implement a kill switch to enact a full shutdown the event of an emergency. The recent White House Executive Order on AI directs federal agencies to determine when and how frontier models may pose national security implications, including developing “know your customer” expectations and safety practices. **SB 1047** creates similar but different regulatory standards for these models. Absent alignment, there could be catastrophic implications for the technology industry in California and the US’s leadership in cloud computing.

- **SB 1047 establishes a new regulatory body with an ambiguous and ambitious purview.** The new “Frontier Model Division” within the Department of Technology would be responsible for a sweeping array of AI-related regulation, including developing novel safety tests and benchmarks, which could very well result in greater inconsistencies with federal rules. Conformity with national and international standards, such as NIST and ISO, should hold authority over those determined by the proposed Frontier Model Division. For example, best practices around red teaming and testing of these covered models are actively being determined by these organizations. Furthermore, additional details and assurances are needed regarding how information and disclosures provided to the Frontier Model Division would be transmitted and stored with the utmost security. Requiring developers and deployers to maintain documentation internally rather than California retaining sensitive, proprietary information on file, would be much more secure. Without clear, realistic requirements, and extraordinary protection of sensitive customer data and proprietary information, developers of frontier AI models are likely to move their training activities and other operations outside of California.
- **SB 1047 imputes excessively harsh penalties, including potentially criminal liability and model deletion.** For instance, developers are required to submit certification of positive safety determinations to the new Frontier Model Division under penalty of perjury, yet the certainty required for that assessment is impracticable if not impossible to obtain. Potential civil penalties include model deletion (in the face of imminent risk or threat to public safety) and “an amount not exceeding 10 percent of the cost, excluding labor cost, to develop the covered model for a first violation and in an amount not exceeding 30 percent of the cost, excluding labor cost, to develop the covered model for any subsequent violation.” Considering the significant resources to train covered models, this sum could amount to many millions.

Ultimately, certain problems demand federal solutions: SB 1047’s inconsistencies will only further fracture the AI regulatory landscape and undermine federal efforts

We cannot overemphasize the importance of ensuring consistency in the AI regulatory landscape, nationally, and the need to follow federal guidance on certain issues that transcend national borders. Relevant to this bill, in October 2023, the White House issued an Executive Order (EO)¹ that requires companies that are developing any foundation model that poses a serious risk to national security, national economic security, or national public health and safety to notify the federal government when training the

¹ [FACT SHEET: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence | The White House.](#)

model and share the results of all red-team safety tests to ensure that AI systems are safe, security and trustworthy before companies make them public.

While we appreciate that in some respects, **SB 1047** appears in line with the goals of the federal government and the White House's EO, the National Institute of Standards and Technology (NIST) is already working with other agencies at the federal level to establish testing and safety guidelines for large models. If enacted, **SB 1047** would likely result in confusion about the correct standards to apply and place additional burdens on AI developers without commensurate gains in safety, especially as it fails to align with regulations nationally and introduces novel concepts and standards including around the assessment of what is a "hazardous capability". Indeed, given the definition of "covered models" under this bill which also scopes in any fine-tuning by downstream customers and users, **SB 1047** is more far-reaching than anything seen to date in voluntary commitments, federal guidance, or the laws of any other countries.

Ultimately, enacting a patchwork of inconsistent AI regulations that go into as much detail as **SB 1047**, will further fracture the U.S. regulatory landscape. As a result, instead of enhancing AI safety, this bill is bound to undermine sensible federal efforts that are already underway and hamper AI innovation in California unnecessarily, encouraging developers to move into other states. Again, this is a conversation that should be had and *is* being had at the national level and there is no need to replicate or duplicate those efforts, particularly in such an inconsistent manner. To the extent that a goal of **SB 1047** might be to set the prevailing standards and practices that the rest of the nation will follow, the lack of clarity and specificity in key definitions outlined above, will only discourage any widespread adoption.

Other recent SB 1047 amendments are more substantive in nature but fail to address our concerns

While the shift from positive safety determination to limited duty exemption is effectively a distinction without a difference, as discussed above, other changes have been made to **SB 1047** that are more substantive nature. Among other things, the amendments include the following:

- The Attorney General is no longer required to commence a civil action when it has "reasonable cause to believe" that a violation has occurred. Instead, the Attorney General is given the discretion to do so, upon finding that a violation has occurred. At the same time, however, the bill, now expressly authorizes punitive damages to be awarded, in addition to other monetary damages and the possibility of an order for the full shutdown of the model as well as other preventative relief that includes deletion of a model and the weights utilized in that model.
- Before initiating training of a frontier model that is not the subject of a limited exemption, and until the model becomes the subject of a limited duty exemption, a developer must now also ensure that their safety and security protocol describes in detail how their testing procedure address the possibility that a covered model can be used to make posttraining modifications or create another covered model in a manner that may generate hazardous capabilities.
- Developers now must provide a reasonable internal process through which an employee can anonymously disclose information to the developer if the employee believes in good faith that the information indicates the developer is out of compliance or has made false or materially misleading statements related to its safety and security protocol. This process includes, at minimum, monthly updates to the employee regarding the status of their disclosure and actions taken in response to the disclosure – presumably in perpetuity, even if the specific issue has been fully addressed, as the bill does not provide any guidance on when those mandated monthly updates can end.
- The new Frontier Model Division must issue guidance on or before July 1, 2026 regarding both the technical thresholds relevant to determining whether an AI model is a covered model, and the technical thresholds and benchmarks relevant to determining if a covered model is subject to a limited duty exemption, taking into account the quantity of computing power used to train covered models that have been identified as having hazardous capabilities and "similar thresholds" used in federal law or regulation for the management of hazardous capabilities. Such guidance is to be updated at least every 24 months after initiate publication.

Unfortunately, none of these changes mitigate concerns we have raised. In some cases, such as in the case of amendments authorizing even harsher penalties (punitive damage) or expanding the already-

sweeping array of AI-related regulations that the new Frontier Model Division would be made responsible for, the amendments in fact exacerbate our stated concerns.

Again, we applaud the intent of this bill but are concerned that its execution will have counterproductive impacts, not only chilling AI innovation, but also preventing AI's beneficial uses. During an incredibly challenging budget year, this bill will result in significant costs to the State in the realm of tens of millions of dollars. In addition to the cost of standing up the new Frontier Model Division and CalCompute, there is also the bigger picture of the incredible potential for future tax revenue that the AI ecosystem can bring to California – meaning, not simply from AI companies, but also from all the industries and businesses looking to leverage AI to increase their efficiency and profitability. Enacting legislation that regulates the development of technology itself, instead of the implementation and uses of it, will be seen as creating a hostile environment for innovation and drive investment to other tech hubs both inside and outside the U.S., with far reaching implications for state revenues. As such, we must unfortunately **OPPOSE SB 1047 (Wiener)**.

Sincerely,



Ronak Daylami
Policy Advocate
on behalf of

Association of National Advertisers (ANA), Christopher Oswald
California Chamber of Commerce, Ronak Daylami
California Land Title Association, Anthony Helton
California Manufacturers and Technology Association (CMTA), Robert Spiegel
Civil Justice Association of California (CJAC), Jaime R. Huff
Computer and Communications Industry Association (CCIA), Naomi Padron
Insights Association, Howard Fienberg
Silicon Valley Leadership Group, Peter Leroe-Muñoz
Software and Information Industry Association (SIIA), Anton van Seventer
TechNet, Dylan Hoffman

cc: Legislative Affairs, Office of the Governor
Severiano Christian, Office of Senator Wiener
Consultant, Senate Appropriations Committee
Ted Morley, Consultant, Senate Republican Caucus

RD:ldl