POLITICOPRO

As Congress stalls on AI regulation, sweeping California proposal would set benchmark

By Eric He

03/14/2024 08:00 AM EDT

Of the many proposals introduced by California lawmakers to regulate artificial intelligence this year, perhaps the most sweeping one would implement a string of regulations for large AI models and set the tone for similar legislation around the country to address the burgeoning technology.

SB 1047, by state Sen. Scott Wiener (D-San Francisco), would require developers of large AI systems to make sure their models cannot perpetuate certain serious public harms — like cyberattacks or the creation of a chemical weapon. Although the models the measure would apply to currently do not exist, they could be developed within the next year.

With Congress stalled on federal legislation around AI, Wiener's bill could serve as a key benchmark for the industry. The proposal separately includes the development of a public cloud for AI research.

At least one prominent tech group — the Chamber of Progress — has voiced opposition, though Wiener told POLITICO he hopes to work together with the industry to find the right balance with the legislation. Wiener said the goal is to preempt potential negative impacts of AI before it is too late, which lawmakers failed to do with social media.

"With more powerful models also comes safety risks, cyber security risks, weapons of mass destruction, and other risks that we need to get ahead of and mitigate," he said.

WHAT'S IN THE BILL?

This Pro Bill Analysis is based on the text of the bill as introduced on Feb. 7.

The Safe and Secure Innovation for Frontier Artificial Intelligence Systems Act (Sec. 1) notes in its declarations that the state Legislature has a role to play in ensuring that California recognizes the benefits of artificial intelligence while avoiding the most severe risks, and that it ensures academic researchers and startups have access to AI innovation (Sec. 2).

The bill would add language to the California Business and Professions Code to both regulate large AI systems and create a public cloud for AI research (Sec. 3).

A "covered model" subject to the bill's proposed regulations would include AI models that are:

— Trained using computer power greater than the 10^26 or floating-point operations in 2024, which is not believed to currently exist

— Reasonably expected to perform similarly to a computer power with 10^26 on benchmarks commonly used to quantify the models per industry best practices

— Below the 10^26 threshold but have a similar general capability

The models would be prohibited from producing systems that have "hazardous capabilities," defined as making it much easier to conduct any of the following "critical harms":

— Creating or using a chemical, biological, radiological or nuclear weapon that results in mass casualties

— Causing a cyberattack on critical infrastructure resulting in at least $500 million in damages

— Engaging in criminal activity that causes at least $500 million in damages

— Other threats to public safety and security comparable to the above

“Derivative models” — or systems that are not independently trained, copied from an existing AI model or consist of a combination of an AI model and other software — would be exempt from the regulations. These would likely be smaller models from startups that do not have the same capacity as large companies to scale up their systems.

Before training the model, a developer should determine whether it can pass a positive safety determination, which is defined as reasonably excluding the possibility that a hazardous capability could occur, and submit certification to the state’s Department of Technology.

Models that do not pass a positive safety determination would be subject to more regulations before they can be trained. Developers would need to take certain safety precautions to prevent theft and misuse, create the ability to fully shut down the AI system and implement guidance set by the state and federal governments, industry best practices and standards-setting organizations.

Developers would have to come up with a written safety and security protocol clearly stating and ensuring that their AI model will not create a hazardous capability, including a detailed description of how it would test for that capability. The protocol, which would be reviewed annually, would also describe conditions that would require a full shutdown and how the safety procedures could be modified. If there is an “unreasonable risk” that the AI model could lead to a hazardous capability, the developer would not be able to start training the model.

Once a model that does not have a positive safety determination is trained, the developer would have to test for that determination and send a notice of compliance with a basis and methodology to the Department of Technology, 30 days after making the model public.

Before making the model public, developers would need safeguards against hazardous capabilities and would have to ensure that any harm that occurs can be traced back to the responsible user. There would also be requirements preventing developers of derivative models from creating a critical harm. If the risk of a hazardous capability remains, the model would not be made public.

Developers would have to submit an annual certification of compliance to the Department of Technology outlining any hazardous capabilities the AI model might possess and whether existing protocols are insufficient to prevent harm.

Any safety incidents by the AI model would have to be reported to the state within four days. These would include:

— A sequence of unsafe behavior that was autonomous and not requested by the user

— Misuses, such as theft, malicious use or inadvertent release, of the model

— Technical or administrative control failure, such as methods to modify the model or limit access to a hazardous capability

— Unauthorized use of a hazardous capability

A developer would be considered in violation of the measure if they proceed with releasing a model that has passed a positive safety determination but there is a comparably powerful model in which a risk of harm has been identified.

Anyone operating a computing cluster — or machines connected by a data center network of over 100 gigabits with the capacity to train AI at a capacity of 10^20 or floating-point operations per second — would have to collect a customers’ information if the customer is able to use the cluster to train a model covered under the bill.

Developers of covered models, or larger AI systems — in addition to a computing cluster — would have to make their pricing transparent for others to purchase access to the model. It would also prohibit discrimination or noncompetitive activity in determining pricing and access.

The measure would not allow for a private right of action over alleged violations, and would instead leave enforcement under the purview of the attorney general. The attorney general would be able to bring forward a civil lawsuit, seeking preventative relief or a restraining order. If there is harm or an imminent risk or threat to public safety, a court would be able to order deletion of the system. Other penalties for violating the measure could include:

— Monetary damages to the victim

— A full shutdown of the system

— A fine of 10 percent of the cost to develop the system, excluding labor costs, for the first violation and a 30 percent fine for each subsequent violation

Defendants would be independently liable for penalties, and corporations found to have taken steps to avoid liability and that are structured in a way wherein paying for damages would be difficult would not be treated as corporate entities.

The measure would offer whistleblower protections to employees seeking to disclose to the attorney general information about their company's violation, and developers of covered models would need to inform employees of their right to come forward (Sec. 3).

Next, the bill would create the Frontier Model Division under the state's Department of Technology to broadly oversee the measure's regulations, including reviewing certification reports from developers and advising the attorney general on potential violations (Sec. 4).

Additionally, the division would be tasked with:

— Issuing additional guidance, standards and best practices to prevent unreasonable risks, and accrediting third parties to certify adherence to those guidelines

— Publishing anonymized AI safety incident reports from developers, and establishing a method for developers to share risk management practices for models that have hazardous capabilities

— Issuing guidance related to how AI could cause a state of emergency and how the governor may respond

— Appointing an advisory committee for open-source AI that can issue guidelines to evaluate models, advising the division on tax credits and incentives for smaller AI developers and consulting on policy

— Levying a fee for developers to submit a certificate

— Developing jury instructions for lawsuits related to a violation of hazardous capability regulations

The measure would also create the Frontier Model Division Programs Fund to collect fees that would be appropriated to carry out the provisions of the bill.

The public cloud established by the measure would be called CalCampute although its creation would be contingent on first having the necessary funding in the budget. The state would be allowed to accept private donations, grants and local funding, as well. CalCompute would focus on research into safely and securely deploying large-scale AI models and fostering equitable innovation. (Sec. 5).

CalCompute would consist of a cloud platform it fully owns and hosts, and the expertise to operate it. The Department of Technology would hire consultants to create CalCompute and would be required to submit an annual report to the Legislature.

The consultants' plan would include an analysis of the cloud platform infrastructure ecosystem to inform the scope of CalCompute, along with establishing partnerships to maintain an advanced computing infrastructure and a framework for which projects the cloud would support.

There would also be a process for evaluating:

— The downstream impact of uses of the public cloud

— The current ability to respond to an emergency

— The progress of collegiate technology-related degree programs

— How CalCompute is retaining workers in the tech sector

The bill would contain a severability clause, and would be "liberally construed" to enact its purposes (Secs. 6, 7). It would apply in conjunction with existing regulations (Sec. 8).

WHO ARE THE POWER PLAYERS?

State Sen. **Scott Wiener** (D-San Francisco) introduced the measure, which is sponsored by the **Center for AI Safety Action Fund**, **Encode Justice** and **Economic Security California**. The Center for AI Safety Action had the largest input when drafting the regulations, while Economic Security California worked mostly on the CalCompute proposal and the youth-led coalition Encode Justice is working to advocate and lobby lawmakers.

Nathan Calvin, senior policy counsel for the Center for AI Safety Action Fund, said in a statement that the proposal accomplishes the goal of ensuring the state "recognizes AI's benefits and adopts industry-leading best practices to avoid its most severe risks, while also making sure AI innovations are accessible to academic researchers and startups."

That access is important, according to **Teri Olle**, director of Economic Security California. In creating CalCompute, Olle hopes to maintain the "lore of the scrappy garage startup" and prevent a few large companies from dominating the AI space as the technology grows.

With a public cloud, a startup would not need the blessing of a tech giant to access resources to develop its AI model. Olle told POLITICO that it is the kind of investment that will help provide a counterweight to the "more entrenched and increasingly-concentrated market ecosystem" around tech and AI.

"It's really exciting to be thinking about a publicly-owned and operated cloud computing cluster that holds open space for development of cutting edge research and innovation that is aligned with the public good — and in service of the public good," Olle said.

The sponsors aimed to create a proposal that covered the biggest harms that can result from AI "in a politically reasonable way," said **Sunny Gandhi**, vice president of political affairs for Encode Justice. Other provisions — like addressing election interference from AI and allowing individuals to sue for violations — were discussed, but ultimately the proposal wound up targeting the most serious harms and setting a higher bar for regulation. Going after AI models that don't yet exist may also limit more vociferous opposition than if current models would be subject to regulation.

Gandhi told POLITICO that Encode Justice — a coalition of 1,000 high school and college students — has a "convincing moral high ground" on AI issues, similar to the climate movement. He said it will be hard for lawmakers to ignore youth advocates and maintain that the harms presented in the bill will not happen if the technology is not regulated.

"We are inhabiting and inheriting the world that these companies are pushing on us, and so we deserve to have our voices heard," Gandhi said. "And I think it does make for a really strong personal argument to all these offices, to see people getting mobilized about this and caring about these kinds of issues."

While Wiener emphasized the benefits of AI and has said he's engaging with stakeholders across the industry, he should expect pushback from powerful tech groups who have deep pockets and have already deployed lobbyists in Sacramento to protect their interests.

Tech groups will be monitoring more than a dozen AI proposals from lawmakers this year, but Wiener's bill, with its strict guardrails, has the potential to cause the most friction with industry groups such as **Chamber of Progress** — a trade industry group founded by a former Google executive whose funders include Amazon, Apple, Cruise and Waymo. The group swiftly condemned Wiener's bill, calling it a "blow to competition."

Todd O'Boyle, senior director of technology policy at Chamber of Progress said such "differential treatment" puts an unfair burden on startups. "Our concern is that this bill, however well-intentioned, is going to limit the equitable distribution of AI models of technology and innovation," he told POLITICO.

Meanwhile, **TechNet**, a trade organization whose members include Meta, Google, Apple and Amazon, said AI has many benefits, but acknowledged that recognizing its risks is crucial.

"America must set the standards for the responsible development and deployment of AI for the world," **Dylan Hoffman**, TechNet's executive director for California and the Southwest, said in a statement. "We look forward to reviewing the legislation and working with Senator Wiener to ensure any AI policies benefit all Californians, address any risks, and strengthen our global competitiveness."

WHAT'S HAPPENED SO FAR?

The California proposal broadly mirrors the Biden administration's Executive Order on AI released last October, which also called for reporting requirements for models trained using a quantity of computing power greater than 10^26 or floating point operations.

But implementing the order has faced an opposition campaign and federal legislation that could codify the proposal has stalled , making it increasingly unlikely that Congress will pass anything before the California Statehouse takes up Wiener's bill.

New York Gov. Kathy Hochul has also proposed a public cloud for AI research in her state using $275 million in state funding and $125 million from colleges. Some experts have called for such a model at the federal level, as well, imagining a public supercomputer that helps agencies solve problems.

WHAT'S NEXT?

SB 1047 is one of more than a dozen AI proposals introduced in the Legislature this year — including three more brought forward on Mar. 13 dealing with AI-generated misinformation ahead of the November election.

Wiener's bill will be heard in the Senate Judiciary Committee on April 2. It has also been referred to the Committee on Governmental Organization.

There could be amendments to the bill depending on discussions and negotiations with industry stakeholders. Julie Rubash, chief privacy officer at the privacy software company Sourcepoint, is watching for how AI will impact privacy and told POLITICO that companies may want to work with lawmakers to craft workable regulations that put everyone on an equal playing field.

"I trust most companies, but it only takes one or two bad actors to cause a lot of harm," Rubash said. "I think that the legislators have recognized that. It's important to address the possibilities of what could happen at the onset, rather than waiting for something catastrophic to happen and then addressing it retroactively."

WHAT ARE SOME STORIES ON THE BILL?

Read POLITICO news on SB 1047.

Lara Korte contributed to this report.