

August 16, 2024

RE: SB 1047 (Wiener) Safe and Secure Innovation for Frontier Artificial Intelligence Systems Act - SUPPORT

Dear Governor Newsom, Senate President pro Tempore McGuire, and Assembly Speaker Rivas,

Notion is a San Francisco-based AI-powered productivity platform serving over 100 million users worldwide. We've witnessed the breathtaking advancement of artificial intelligence firsthand. Every day, we design products that use AI foundation models to enhance creativity and productivity for individuals and businesses across the globe. It's clear to me that these systems will become dramatically more powerful over the next few years, unlocking new realms of innovation in science, medicine, and beyond.

However, as AI systems grow more sophisticated, we must grapple with the very real risks they pose. Without reasonable precautions, AI could cause severe harms on an unprecedented scale - from cyberattacks on critical infrastructure to the development of chemical and biological weapons to automated cybercrime. These risks are not distant science fiction scenarios; they are urgent challenges that demand our immediate attention.

SB 1047: A Prudent Balance

California's SB 1047 strikes a prudent balance between fostering innovation and protecting public safety. It focuses on common-sense regulation for the handful of companies developing the most powerful AI systems, while preserving the innovative spirit that has made California the global leader in technology.

A key virtue of SB 1047 is that its requirements align closely with industry best practices that many responsible companies are already implementing. Good corporate citizens deserve recognition, but voluntary measures are not sufficient to protect public safety. SB 1047 would require companies developing the most powerful AI systems to test for safety risks before releasing them. If testing reveals that an AI system could potentially cause catastrophic harm, the company must implement reasonable safeguards to mitigate these risks.

Critically, the bill includes robust whistleblower protections for employees who report safety concerns at AI companies. These protections are essential because these employees will be the first to know if a company's actions put the public at risk.

Regulation should focus on foundation models, just not specific uses

Some critics argue that regulation should focus solely on harmful uses of AI, rather than the underlying technology. This view is dangerously misguided. It's already illegal to use AI to commit crimes, so adding a duplicative law won't make Californians any safer. And relying on laws that make it illegal to cause these harms isn't sufficient to keep Californians safe; arguing otherwise is like claiming we don't need background checks on guns because assault and murder are already illegal. The way to prevent extremely dangerous capabilities from falling into the wrong hands is to ensure appropriate safeguards are in place at the foundation model level.

SB 1047 wisely focuses on regulating the small number of well resourced companies building the biggest AI systems, which cost hundreds of millions of dollars to develop. It does not regulate startup companies at the heart of California's innovation economy. As a company relying on these models, we believe this approach is the most effective way to mitigate risks across the entire AI ecosystem without burdening smaller developers and startups.

Our software builds on AI foundation models, testing and safety practices allow us to integrate AI into our products with the confidence that doing so is safe and reliable for our tens of millions of users. Some critics claim that SB 1047 would hinder startups, but the opposite is true. If an AI product were to cause a critical harm, this could have disastrous effects for consumer trust in AI and create ripples throughout the ecosystem. It would be an unrealistic and crippling burden to exclusively task application developers with ensuring the safety of AI systems they didn't develop. Far from banning open-source models and the startups that use them, SB 1047 would facilitate justified trust in the products that founders build upon every day. Incentivizing robust and reliable frontier models at the foundation level is the most effective way to ensure safety throughout the AI value chain without burdening small developers and startups.

SB 1047 is an opportunity to lead on responsible innovation

Some have expressed concerns that this bill might stifle innovation or cause companies to leave California. Notion strongly disagrees. The measures required by SB 1047 are not a burden but a necessary investment in building trust and ensuring our technology is developed responsibly. California's unique ecosystem of talent, capital, and innovation makes it the ideal place to develop AI responsibly. By taking the lead on responsible AI development, California will strengthen its position as a global technology leader.

Ultimately, SB 1047 represents an opportunity to be seized. Tech leaders from Meta's Mark Zuckerberg to OpenAI's Sam Altman have appeared before Congress to discuss AI regulation, warn of potentially catastrophic effects, and even request regulatory guidance. However, expectations for federal action remain low. California has the chance to lead the way in establishing a framework for responsible AI development.

At Notion, we would not face any new burdens under SB 1047 and we are not concerned about indirect effects on our work due to regulations on the companies we partner with. On the contrary, we believe that reasonable regulation like SB 1047 will prevent the kinds of AI safety incidents that could result in knee-jerk overregulation and stifle innovation in the future.

Criticisms of SB 1047 mostly miss the mark

We've seen criticism of SB 1047 that misses the mark. One such argument is that SB 1047 requires developers to "predict every possible use of their model" and demonstrate that misuse is impossible. However, SB 1047 only requires providing *reasonable assurance* of preventing *unreasonable risks* of specific *critical harms* that could result in mass casualties or \$500M of damage. If a company spending hundreds of millions to develop AI models (whether open- or closed-source) is unable to provide any such reassurances, they clearly should not be allowed to recklessly release the model they believe poses an unreasonable risk to public safety.

Another argument is that SB 1047's *full shutdown* requirement is impossible to comply with for open source developers. However, this rests on an inaccurate reading of the bill. The *full shutdown* requirement only affects models controlled by a developer. Once a model has been open-sourced, the developer is not required to be able to shut down the copies that are no longer in its possession. Opponents claiming otherwise are misleading their readers.

Others have claimed that the 10^{26} floating point operation threshold for computing power on a single training run is arbitrary, and that a FLOP threshold is unreasonable. But a core lesson of the past decade of AI progress is that model capabilities are strongly predicted by the amount of computing power used to train those models. Today's models do not possess these hazardous capabilities, but the most powerful models are already showing some warning signs. As a result, it would be prudent to conduct capabilities evaluations for the next generation of models. No models released today have been trained on over 10^{26} FLOP, but the next generation of models will very likely cross this threshold. This is why President Biden's Executive Order on AI used the very same threshold that SB 1047 proposes.

A Foundation for the Future

By making smart policy choices now, we can create the foundation for a thriving, competitive AI industry that improves our quality of life while keeping Californians safe. SB 1047 provides a clear framework for compliance that offers developers a roadmap for responsible innovation today while allowing room for adaptation in this rapidly evolving field.

We urge you to support and swiftly pass this crucial legislation. The rapid pace of AI development makes it imperative that these guardrails are put in place as soon as possible. Notion is committed to working with policymakers, fellow developers, and the public to ensure AI benefits society, and we stand ready to provide any further information or assistance you may need in considering this bill.

Thank you for your consideration.

Sincerely,

Simon Last
Cofounder, Notion