June 11, 2024

The Honorable Rebecca Bauer-Kahan
Capitol Office
1021 O Street
Suite 6320
Sacramento, CA 942849-0001

Dear Chair Bauer-Kahan,

BSA | The Software Alliance appreciates the opportunity to share insights from the enterprise software sector on artificial intelligence (AI) and SB 1047, the Safe and Secure Innovation for Frontier Artificial Intelligence Models Act. BSA is the leading advocate for the global software industry.[1] BSA members are at the forefront of developing innovative services, and their products are used by businesses of all sizes across every sector of the economy. AI is much more than robots, self-driving vehicles, or social media; it is used by companies large and small to create and improve the products and services they provide to consumers, to streamline their internal operations, and to enhance their capacity to make data-informed decisions. BSA members are on the leading edge of providing businesses-to-business tools that help companies leverage the remarkable benefits of AI.[2]

As leaders in the development of enterprise AI, BSA members have unique insights into the technology's tremendous potential to further spur digital transformation in the private and public sectors and the policies that can best support the responsible use of AI, especially high-risk AI. BSA's views are informed by our experience with members developing the BSA Framework to Build Trust in AI,[3] a risk management framework for mitigating the potential for unintended bias throughout an AI system's lifecycle. Built on a vast body of research and informed by the experience of leading AI developers, the BSA Framework outlines a lifecycle-based approach for performing impact assessments to identify risks of AI bias and highlights corresponding risk

[1] BSA's members include: Adobe, Alteryx, Asana, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, Cohere, Databricks, DocuSign, Dropbox, Elastic, Graphisoft, Hubspot, IBM, Informatica, Kyndryl, MathWorks, Microsoft, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, Prokon, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Splunk, Trend Micro, Trimble Solutions Corporation, TriNet, Twilio, Workday, Zendesk, and Zoom Video Communications, Inc.

[2] *See* BSA | The Software Alliance, Artificial Intelligence in Every Sector, *available at* https://www.bsa.org/files/policy-filings/06132022bsaaieverysector.pdf.

[3] *See* BSA | The Software Alliance, Confronting Bias: BSA's Framework to Build Trust in AI, *available at* https://www.bsa.org/reports/confronting-bias-bsas-framework-to-build-trust-in-ai.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

mitigation best practices. BSA's extensive experience has helped us identify effective policy solutions for addressing AI risks.

We appreciate the intent of SB 1047 and share the goal of promoting the responsible and safe development and use of AI. We believe governments can play a vital role in establishing guardrails around the creation and deployment of AI technologies. However, we are concerned that SB 1047 establishes an extremely prescriptive and complex regulatory regime that risks stifling, if not eliminating, the development of frontier models in California—even as those models continue to be developed elsewhere. The bill also upends the roles and responsibilities of different actors in the AI ecosystem in ways that do not improve consumer safety or security.

As currently drafted, BSA is concerned that the bill:

- Broadly regulates a single technology—frontier models—without creating risk-based policies aimed at high-risk uses of AI;
- Establishes an extremely prescriptive, complex regulatory regime;
- Creates infeasible exceptions;
- Blurs the roles and responsibilities of different actors in the AI value chain, including through new pre-training obligations and pre-market obligations;
- Creates vague and unworkable safety incident reporting obligations; and
- Establishes a new government agency with a broad and vague mandate, as part of its unclear enforcement and compliance mechanisms.

I. The Bill Should Focus on Risks, Rather Than Broadly Regulating Frontier Models

Instead of creating legislation that addresses high-risk uses of AI, SB 1047 broadly regulates developers of frontier AI models. Specifically, SB 1047 focuses on "covered models," meaning AI models that: (1) were trained using computing power greater than 10^26 integer or floating-point operations in 2024; and (2) cost over $100 million to create.

This approach ignores the context in which AI models are deployed. For example, covered models may support companies across a wide range of industries, such as critical infrastructure and health care. They can also help individuals create new recipes or write emails. Each use will present different safety and security risks. Broadly regulating covered models fails to recognize the range of unique circumstances in which AI models are deployed—including both high-risk and low-risk uses—and ignores that different risk mitigation approaches may be appropriate in different scenarios.

More broadly, SB 1047's focus on frontier models also risks creating conflicts with international efforts aimed at supporting the development and deployment of safe and trustworthy AI worldwide. For example, in the United States, the federal government now requires companies developing frontier models to report certain metrics to the Department of Commerce, under an Executive Order issued last October. The Group of Seven countries have also developed a code of conduct for advanced AI systems, to support an internationally harmonized approach to the responsible development of new AI systems. In contrast, SB 1047 would empower a new state-level agency to enforce extremely prescriptive and complex regulatory requirements. Instead, we recommend

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001
P 202-872-5500
W bsa.org

that SB 1047 encourage companies to adopt effective practices they can undertake now to promote the safety and security of AI technologies, rather than inadvertently hampering developers' ability to train covered models.

II. The Bill Creates an Extremely Prescriptive, Complex Regulatory Regime

The bill creates an extremely prescriptive and complex regulatory regime for companies that develop frontier AI models. Under the bill, developers of such models must meet certain obligations before they train the model, after they train the model, and before the model is offered to the public. At each turn, these obligations are tied to vague concepts of harm (including harm that may be caused by the developer's own model or by other models created by other actors) and yet-to-be written guidance, while requiring extensive disclosures to a new state agency. For example:

- *Before starting to train a model*, the developer must adopt a list of safety and security measures that are tied to forthcoming guidance to be issued from the new state agency and guarantee that *other actors* won't use their model to create another model that may cause certain "hazardous" results.

- *Before public use of the model*, the developer must implement safeguards for both its own models and for future models derived from its model; it must also implement "other measures reasonably necessary" to protect against harms, a mandate that is vaguely tied to future guidance from the new state agency, as well as from the National Institute of Standards and Technology and standard-setting organizations.

Developers must also submit a host of materials about their compliance with these requirements to the new state agency, under penalty of perjury.

These requirements far exceed requirements of the EU AI Act, which creates obligations for companies that develop general purpose AI models with "systemic" risks. Models are presumed to have such risks if the cumulative amount of compute used to train the model is greater than 10^25 floating point operations—a threshold similar to the threshold for "covered models" under SB 1047. In the EU, providers of GPAI models with systemic risks are subject to heightened obligations under Article 55 of the EU AI Act. Those obligations include requirements to evaluate the model, assess and mitigate possible systemic risks, document serious incidents, and ensure appropriate cybersecurity protections.

III. The Bill's Limited Duty Exemption is Unworkable

In theory, developers can be exempt from SB 1047's regulatory requirements in certain circumstances—but the bill's exceptions are not workable in practice, even after recent amendments.

Under the bill, a developer may determine before training a covered model if that model qualifies for a "limited duty exemption" from the bill's obligations. This exception applies if the developer can provide reasonable assurance that the covered model does not have a hazardous capability and will not come close to possessing a hazardous capability when accounting for a reasonable margin for safety and the possibility of posttraining modifications.

It is unclear that any covered model could qualify for this exemption—because it relies on an ambiguous definition of hazardous capability. Under the bill, hazardous capability means the ability of a covered model to "enable" certain harms in a way that would be "significantly more difficult to cause without access" to the covered model. The bill identifies harms including

chemical and biological weapons, $500 million of damage in cyberattacks, and $500 million of damage by criminal activities. This definition is met even if the hazardous capability would not occur but for fine tuning and posttraining modifications performed by third-party experts intended to demonstrate those abilities. The definition is unclear in at least three ways:

- It is unclear how a covered model "enables" harms, or how a developer would measure whether the model does so "in a way that would be significantly more difficult to cause" without access to the model.

- The definition relies on damage thresholds for several harms; these are difficult to measure against.

- Including harms that arise through fine tuning and posttraining modification requires the developer to represent what other actors can do, which conflates their different roles.

Although recent amendments appear to encourage use of this exception—by ensuring it requires a developer provide only "reasonable assurance" against such harms—it is unclear that developers may use this exemption in practice.

IV. The Bill Blurs the Roles of Developers and Deployers

BSA members are at the forefront of developing best practices for AI safety and security. We believe governments can play a key role in encouraging companies to build safeguards around their creation and use of AI technologies. However, SB 1047 places a concerning set of obligations on covered model developers that may discourage thorough training and be counterproductive to promoting safety and security. These requirements also blur the roles of developers and deployers, in both the pre-training obligations and pre-market obligations.

a. Pre-Training Obligations

We provide the following feedback on Sec. 22603(b)'s obligations for developers before initiating training of a covered model:

First, several obligations hold covered model developers responsible for downstream uses. The bill requires developers of covered models to meet a host of obligations before they train the covered model. However, these obligations blur the distinction between AI developers and AI deployers by holding developers of a covered model responsible for actions of other companies that may later develop derivative versions of that model.

The developer of a covered model is not situated to assess the risks of later derivative models. That is because determining whether a derivative model potentially has hazardous capabilities will largely depend on the context in which the derivative model is deployed and mitigation measures undertaken by the deployer, which is often a different entity than the developer of the model. Put simply, these requirements hold covered model developers responsible for downstream uses over which they have no control.

For example, the bill requires the developer of a covered model to:

- Implement a safety and security protocol that manages the risks of developing and using covered models across their life cycle, including risks posed by "enabling or potentially

enabling the creation of derivative models." Under the bill, the protocol is to provide reasonable assurances that the developer will not enable the production of a derivative model with a hazardous capability.

- Refrain from initiating the training of a covered model if there remains an unreasonable risk that an individual, or the covered model itself, may be able to use the hazardous capabilities of the covered model, or a derivative model based on it, to cause a critical harm.

- Implement the capacity to promptly enact a full shutdown of the covered model, including a shutdown of derivative models.

Second, we recommend maintaining recent changes to the definition of full shutdown that consider open-source development. Recent amendments to the bill have improved its application to open-source technologies. Open source is a critical component of the AI ecosystem. It expands the AI marketplace, enhances the diversity of product offerings, promotes transparency, and enables vulnerabilities to be identified and remediated. Because open-source models can be freely downloaded and modified, open-source model developers have no means to restrict or influence use of the open-source model.

We agree with the amendments to the definition of a full shutdown, which reflect these issues. Under the amendments, where the bill requires certain covered model developers to implement the capacity to promptly enact a full shutdown, we appreciate that the definition of full shutdown now accounts for open-source development.

b. Pre-Market Obligations

The bill also places a concerning set of pre-market obligations on certain covered model developers in Sec. 22603(d). Again, these obligations blur the distinction between AI developers and AI deployers.

For example, the bill requires the developer of a covered model to:

- Implement reasonable safeguards and requirements, informed by the training and testing process, to prevent an individual from using the hazardous capabilities of the covered model or a derivative model to cause a critical harm or from using the model to create a derivative model that is used to cause a critical harm;

- Provide reasonable requirements to developers of derivative models to prevent the derivative model's use to cause a critical harm; and

- Refrain from making the covered model widely available if there "remains an unreasonable risk" that the hazardous capabilities of the model or a derivative model could be used to cause a critical harm.

Similar to several pre-training activities required in Sec. 22603(b), these requirements hold covered model developers responsible for downstream uses over which they have no control.

V. The Bill's Safety Incident Reporting Requirements are Vague and Unworkable

The bill imposes vague and unworkable safety reporting requirements on covered model

developers. Sec. 22603(g) requires covered model developers to report each "artificial intelligence safety incident" to the new regulatory agency. Reports must be made no later than 72 hours after learning an AI safety incident has occurred or is reasonably likely to have occurred. This requirement creates two significant concerns:

First, AI safety incident is broadly and vaguely defined. The bill defines an AI safety incident as including a broad range of events, such as the covered model autonomously engaging in behavior that materially increases the risk of a hazardous capacity being used; theft, misappropriation, inadvertent release, unauthorized access, or escape of the model weights; and critical failure of technical or administrative controls. This definition is broad or vague in at least three ways:

- Hazardous capability is broadly and vaguely defined, as described above.
- The definition of AI safety incident captures incidental circumstances that may not have any effect on the safety of the covered model. For example, if, due to a clerical error, a company employee temporarily gains access to a covered model's model weights but makes no changes, this "unauthorized access" may have no effect on the covered model's outputs. However, the bill would consider this accidental and inconsequential event reportable.
- "Critical failure" is undefined.

Given the ambiguity surrounding the definition of AI safety incident, we are concerned this requirement would result in over-notification to the new regulatory agency created by the bill and thereby not promote safety.

Second, the timeline to report AI safety incidents is unreasonable. The 72-hour timeline for reporting AI safety incidents creates significant practical challenges, because companies will often lack a comprehensive understanding of an incident within that limited time frame. Further, the notification requirement obligates companies to divert resources from addressing the safety of systems to fulfilling short-fuse reporting requirements, as an incident is unfolding.

VI. Establishing a New Government Agency with A Broad and Vague Mandate Results in Unclear, Vague Compliance

The bill's complicated, prescriptive requirements are to be enforced by a new agency—which is given broad, vague, and ambiguous authority. The bill also lacks clear and consistent enforcement and compliance mechanisms, which risk creating confusion for businesses seeking to understand their obligations.

These provisions raise at least four distinct concerns:

First, the bill grants significant enforcement authority to a newly created regulatory agency. The bill's creation of a new regulatory body in the Frontier Model Division within the Department of Technology is concerning, particularly considering the new regulator's broad authority to issue guidance, standards, and best practices that will become binding under the bill. Existing technical standards for AI are nascent and should be developed consistent with longstanding voluntary, market-driven, and consensus-based approaches to standards development. Tasking a new regulatory with creating binding guidance in a quickly evolving area is unlikely to create robust and predictable guidance for companies.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

Second, the Department of Technology has little experience regulating the private sector. It is unclear why the bill creates this new authority within the Department of Technology, which is generally focused on the state government's use of technology and the delivery of digital government services. Creating a new division within the agency to regulate commercial uses of AI would be a significant departure from the department's existing work.

Third, the bill requires compliance under penalty of perjury. The bill requires covered model developers to provide annual certifications to the new regulator established by the bill *under the penalty of perjury*. Those certifications must be signed by the developer's chief technology officer or a more senior corporate officer. This approach subjects corporate officers of covered model developers to potential criminal liability for certifying to the bill's ambiguous requirements. Such penalties are disproportionate to the harms that SB 1047 seeks to address.

Fourth, the requirements for compliance certifications undermine the bill's goal of increasing safety. Covered model developers completing the bill's required compliance certifications must include specific information about the covered model, including an assessment of how the safety and security protocol may be insufficient to prevent harms from the covered model's hazardous capabilities. Requiring covered model developers to provide this information is unreasonable, because providing an assessment of the deficiencies of the safety and security protocol may inadvertently provide malicious actors with instructions for how to exploit the covered model.

* * *

Thank you for allowing us to provide the enterprise software sector's perspective. We welcome the opportunity to serve as a resource and further engage with you or a member of your staff on these important issues.

Sincerely,

Meghan Pensyl

Meghan Pensyl
Director, Policy

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org