August 21, 2025

The Honorable Buffy Wicks
Chair, Assembly Committee on Appropriations
1021 O Street, Suite 8220
Sacramento, California 95814

Re: Senate Bill 53 – AI Models: Large Developers - OPPOSE

Dear Chair Wicks, Vice Chair Sanchez, and Members of the Committee on Appropriations,

The Business Software Alliance appreciates the opportunity to share insights from the enterprise software sector on artificial intelligence (AI) and Senate Bill 53, the Transparency in Frontier Artificial Intelligence Act, which we **oppose**. BSA is the leading advocate for the global software industry.[1] BSA members are at the forefront of developing cutting-edge services, and their products are used by businesses of all sizes across every sector of the economy.[2]

AI is changing the way we live and work, and it has real-world benefits. Realizing the potential of AI requires trusting that the technology is developed and deployed responsibly.
Crafting AI legislation that promotes responsible uses of AI and protects against misuse is one of the most important technology issues today, and one we already see governments beginning to tackle, including in the European Union and in Colorado. The most effective way to address this issue is through a single, national law. However, just as states took the lead in adopting consumer privacy laws, we recognize states are again leading with AI legislation.

We appreciate the intent of SB 53 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 53 establishes an extremely prescriptive and complex regulatory regime that risks stifling the development of frontier models in California—even as those models continue to be developed elsewhere.
Specifically, BSA is concerned that SB 53:

[1] BSA's members include: Adobe, Alteryx, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cohere, Dassault Systemes, Databricks, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Informatica, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Trend Micro, TriNet, Workday, Zendesk, and Zoom Communications Inc.
[2] *See* The Business Software Alliance, Artificial Intelligence in Every Sector, *available at* https://www.bsa.org/files/policy-filings/06132022bsaaieverysector.pdf.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

- Broadly regulates a single technology—foundation models—without creating risk-based policies aimed at high-risk uses of AI;
- Establishes a prescriptive regulatory regime;
- Relies on a conception of catastrophic risk and dangerous capability is unworkable;
- Prematurely relies on third-party audits; and
- Creates vague and unworkable safety incident reporting obligations.

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

Instead of creating legislation that addresses high-risk uses of AI, SB 53 broadly regulates foundation AI models, meaning models that are: (1) trained on a broad data set; (2) designed for generality of output; and (3) adaptable to a wide range of distinctive tasks. The bill applies to large developers of such models, meaning: (1) the person has trained, or initiated the training of, at least one foundation model using a quantity of computing power greater than 10^26 integer or floating-point operations; and (2) the person had annual gross revenues in excess of one hundred million dollars in the preceding calendar year. After January 1, 2027, the Attorney General can update the definition of large developer.

This approach ignores the context in which AI models are deployed. For example, foundation 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 the largest foundation 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.

Furthermore, the bill's one-size-fits-all approach ignores the different risk mitigation capabilities of different users. Enterprise deployers often use foundation models designed for business-to-business contexts and have significant safety and security resources, including in-house experts and governance policies. Such companies have far greater capabilities to safeguard against the safety and security risks associated with their use of a foundation model than individuals using a foundation model for their personal use.

More broadly, SB 53's focus on foundation models also risks creating conflicts with international efforts aimed at supporting the development and deployment of safe and trustworthy AI worldwide. For example, 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 53 would establish prescriptive regulatory requirements that risk inadvertently hampering developers' ability to train foundation models.

II. The Bill Creates a Prescriptive Regulatory Regime.

The bill creates a prescriptive regulatory regime for companies that train large foundation models. Under the bill, large developers must write, implement, and post publicly on their website a safety and security protocol that includes ten detailed pieces of information about the foundation model, including the:

- Testing procedures used to assess catastrophic risks from the foundation model;
- Actions the large developer will take if the foundation model has attained a dangerous capability or poses a catastrophic risk;
- Mitigations taken to reduce catastrophic risk and how the large developer assesses the effectiveness of those mitigations;
- Degree to which the large developer's assessments of catastrophic risk and dangerous capabilities and the effectiveness of catastrophic risk mitigations are reproducible by external entities; and
- Large developer's cybersecurity practices and how the large developer secures unreleased model weights from unauthorized modification or transfer by internal or external parties.

Concerningly, requiring large developers to publicly publish their safety and security protocol may inadvertently create safety and security risks. Particularly because the safety and security protocol must include detailed information about the large developer's risk assessments and mitigations, publicly publishing it risks creating a roadmap for bad actors to exploit.

III. The Bill's Definition of Conception of Catastrophic Risk and Dangerous Capability is Unworkable.

The bill focuses on whether foundation models pose "catastrophic risk" or attained "dangerous capability" and requires large developers to test and create mitigations based on those terms. Both terms, however, create a vague threshold that will cause significant confusion for companies seeking to implement their obligations.

Catastrophic risk is defined as "a foreseeable and material risk that a large developer's development, storage, use, or deployment of a foundation model will materially contribute to the death of, or serious injury to, more than 50 people or more than one billion dollars ($1,000,000,000) in damage to, or loss of, property arising from a single incident, scheme, or course of conduct involving a dangerous capability."

Dangerous capability is defined as "the capability of a foundation model to do any of the following: (1) provide expert-level assistance in the creation or release of a chemical, biological, radiological, or nuclear weapon; (2) conduct or assist in a cyberattack; (3) engage in conduct, with limited human intervention, that would, if committed by a human, constitute the crime of murder, assault, extortion, or theft, including theft by false pretense; (4) evade the control of its large developer or user."

The definition of catastrophic risk's focus on how a foundation model "materially contribute[s]" to the number of deaths or serious injuries and amount of damage is overly broad and may be indeterminable. Similarly, the definition of dangerous capability is overly broad, including by focusing on a foundation model's ability to "assist in" a cyberattack and "evade" control. Together, the definitions do not create a measurable threshold against which large developers can assess a foundation model.

IV. The Bill's Reliance on Third-Party Audits Is Premature.

The bill requires large developers, beginning in January 2030, to retain a third-party auditor to assess the large developer's compliance with its safety and security protocol. Although the effective date of such obligations is delayed by several years, it's unclear whether the AI auditing ecosystem will be sufficiently mature to support these obligations. Specifically, this approach raises at least two distinct concerns:

First, the AI audit ecosystem is still immature and is not fully equipped to assure the performance of an AI system. There are few existing procedures or best practices for businesses to choose a reputable company capable of auditing an AI system or foundation model. It is not clear how an auditor would assess a large developer's compliance with its safety and security protocol given the broad and vague definitions in the bill, as described above. Nor is there a comprehensive, uniform, and appropriate set of standards any such auditing company should apply to assess an AI system more broadly. As in regulated industries like accounting and financial services, professional standards and guidelines for auditing are needed to create a functional and mature AI audit ecosystem.

Second, third-party audits create confidentiality and intellectual property concerns. Third-party audits often require access to confidential business proprietary information, creating concerns around the treatment of trade secrets and other sensitive information. Moreover, in some cases, audits may require access to consumers' personal information,

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

and providing large data sets associated with AI systems to a third-party auditor can create significant privacy concerns. At the same time, if an auditor is provided access to the AI system itself, it can create security concerns. Engaging a third-party auditor would require working through these and other issues that are not presented in the context of internally focused accountability tools, such as impact assessments.

Moreover, the bill requires an auditor to transmit a high-level summary of the outcome of the audit, and offers no protection for companies' trade secrets, intellectual property, or other confidential information. Such a requirement compounds the confidentiality and intellectual property concerns posed by third-party audits.

In light of these challenges and the nascent AI auditing landscape, independent third-party audits are not currently the most effective mechanism for managing AI risks. Instead, organizations should leverage other risk mitigation measures, such as those outlined in the NIST AI RMF, and conduct internal impact assessments to identify and mitigate risks.

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

The bill imposes vague and unworkable safety reporting requirements on large developers. Large developers must report "critical safety incident[s]" pertaining to one of its foundation models to the Attorney General within 15 days of discovery and within 24 hours if the critical safety incident poses an imminent risk of death or serious physical injury. This requirement creates two significant concerns:

First, critical safety incident is broadly and vaguely defined. The bill defines a critical safety incident as meaning "any of the following: (1) Unauthorized access to, modification of, or exfiltration of, the model weights of a foundation model; (2) Harm resulting from the materialization of a catastrophic risk; (3) Loss of control of a foundation model causing death, bodily injury, or damage to, or loss of, property; (4) A foundation model that uses deceptive techniques against the large developer to subvert the controls or monitoring of its large developer outside of the context of an evaluation designed to elicit this behavior; (5) Attaining a dangerous capability or catastrophic risk threshold, as defined in the large developers' safety and security protocol [as required by the bill], for the first time." This definition is broad or vague in at least three ways:

- Dangerous capability and catastrophic risk are 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.

- It's unclear if the "harm" resulting from the materialization of a catastrophic risk is limited to the vague thresholds created by the definition or some other, undefined harm.

Given the ambiguity surrounding the definition of critical safety incident, we are concerned this requirement would result in over-notification and thereby not promote safety.

Second, the timeline to report AI safety incidents is unreasonable. The 15-day timeline for reporting AI safety incidents creates significant practical challenges, including because the requirement hinges on discovery of a critical safety incident, which is vague. Companies often must quickly and thoroughly investigate whether an incident has occurred, and it's unclear if the reporting requirement is triggered when a company first suspects, reasonably believes, or confirms a critical safety incident has occurred. Further, the reporting requirement obligates companies to divert resources from addressing the safety of systems to fulfilling short-fuse reporting requirements, as an incident is unfolding.

* * *

Thank you for allowing us to provide the enterprise software sector's perspective on these issues. 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
Director, Policy

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org