April 1, 2024

The Honorable Thomas J. Umberg
Senate Judiciary Committee
1021 O Street, Room 3240
Sacramento, CA 95814

Dear Chair Umberg,

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

[1] BSA's members include: Adobe, Alteryx, Asana, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, Databricks, DocuSign, Dropbox, Elastic, Graphisoft, Hubspot, IBM, Informatica, Kyndryl, MathWorks, Microsoft, Okta, 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

We appreciate the intent of SB 1047 and share the goal of helping to ensure that AI is developed and used safely and securely. We believe governments can play a vital role in encouraging best practices and promoting guardrails around the creation and deployment of AI technologies.

We are concerned with SB 1047's focus on specific types of AI technology, rather than focusing on high-risk uses of AI. Instead of regulating a specific type of technology, we believe policymakers addressing AI issues should focus on the priorities outlined below. We therefore recommend SB 1047 be revised to: (1) focus on high-risk uses of AI; (2) create important guardrails on those uses, by requiring risk management programs and impact assessments; (3) distinguish between different entities in the AI value chain, including developers and deployers; and (4) adopt a clear and consistent enforcement mechanism.

I. Focus on High-Risk Uses

BSA recommends you and the committee focus on high-risk uses of AI, particularly AI systems that determine an individual's eligibility for housing, employment, credit, education, access to physical places of public accommodation, healthcare, or insurance. These systems have the potential to affect important life opportunities — and are a key area for policymakers to address. In contrast, many everyday uses of AI present few risks to individuals and create significant benefits, like helping organize digital files, auto-populate common forms for later human review, improve a company's ability to forecast supply chain issues, and detect, prevent, and respond to cybersecurity threats.

As currently drafted, SB 1047 broadly regulates developers of certain AI models. Specifically, SB 1047 applies to covered models, meaning AI models that: (1) were trained using computing power greater than 10^26 integer or floating-point operations in 2024; or (2) are reasonably expected to have similar or greater performance. By regulating these models broadly, SB 1047 ignores the context in which these AI models are deployed. For example, covered models could one day be deployed in a wide range of contexts, from critical infrastructure or health care, to helping individuals create new recipes or draft emails. Each of these use cases will present different safety and security risks, and broadly regulating these models fails to recognize how the unique circumstances in which AI models are deployed will influence risk mitigation approaches. We recommend SB 1047 focus on high-risk uses of AI to address the most important impacts AI has on consumers' lives.

II. Encouraging Best Practices for Safety and Security

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 testing and be counterproductive to promoting safety and security.

Specifically, Sec. 22603(a) requires developers of covered models, prior to initiating any training, to assess whether they can make a positive safety determination, meaning the developer can reasonably exclude the possibility that a covered model has hazardous capability or may come close to possessing a hazardous capability, where hazardous capability includes chemical, biological, radiological, or nuclear weapons, incidents of a certain damage threshold, or other similarly severe threats to public safety and security. Positive safety determination and hazardous capability are overly broad, and the bill is unclear as to what constitutes potentially coming "close to" possessing a hazardous capability. The bill also concerningly requires covered

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

model developers to certify to the new regulator established by the legislation *under the penalty of perjury* that it has made a positive safety determination, which could subject developers of covered models to criminal liability for certifying to ambiguous requirements. Imposing criminal penalties under these circumstances would be disproportionate to the harm the requirement seeks to address. Moreover, the type of training this section precludes is the exact type of activity a developer would use to assess the risks posed by the covered model.

Additionally, standards for AI safety and security are rapidly evolving, with groups like the US AI Safety Institute working to support the development and deployment of safe and trustworthy AI. In addition to focusing on high-risk uses of AI, we recommend SB 1047 encourage effective practices companies can undertake now to promote the safety and security of AI technologies, rather than inadvertently hampering developers' ability to train and test covered models.

We recommend the bill adopt a different approach, by focusing on high-risk uses of AI and creating important guardrails on those uses by requiring risk management programs and impact assessments, as described in more detail below.

a. Risk Management Programs

Companies should implement risk management programs that help them identify and mitigate risks. Risk management programs establish repeatable processes for companies to identify and mitigate potential risks that can arise throughout the lifecycle of an AI system. Risk management is particularly important in contexts like AI, privacy, and cybersecurity, where the combination of quickly evolving technologies and highly dynamic threat landscapes can render traditional approaches to compliance ineffective. Risk management programs have two key components: (1) a governance framework of policies, procedures, and personnel that support the company's risk management function, and (2) a scalable process for performing impact assessments that identify and mitigate risks of an AI system.

One way for companies to establish risk management programs is by using the AI Risk Management Framework (AI RMF), which was released last year by the National Institute of Standards and Technology (NIST).[4] The AI RMF builds on NIST's work creating frameworks for managing cybersecurity and privacy risks.[5] The AI RMF helps companies incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products. Ultimately, effective AI risk management programs should support coordination across the company, to promote the identification and mitigation of risks throughout the lifecycle of an AI system.

b. Impact Assessments

BSA recognizes that performing impact assessments for high-risk uses of AI is a key part of creating a meaningful risk management program. Impact assessments have three purposes: (1) identifying potential risks that an AI system may pose, (2) quantifying the degree of potential harms the system could generate, and (3) documenting steps taken to mitigate those risks.[6] Impact assessments are already widely used in a range of other fields, including privacy, as an accountability mechanism that demonstrates a product or system has been designed in a manner

[4] *See* NIST AI Risk Management Framework, *available at* https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.
[5] *See* NIST, Cybersecurity Framework, Questions and Answers, (discussing federal agency use of the NIST CSF), *available at* https://www.nist.gov/cyberframework/frequently-asked-questions/framework-basics#agency.
[6] See BSA, Impact Assessments: A Key Part of AI Accountability, *available at* https://www.bsa.org/files/policyfilings/08012023impactassess.pdf.

that accounts for the potential risks it may pose to the public. Because impact assessments already exist today, they can be readily adapted to help companies identify and mitigate AI-related risks.[7] In our view, when AI is used in ways that could adversely impact civil rights or access to important life opportunities, the public should be assured that such systems have been thoroughly vetted and will be continuously monitored to account for the risks associated with unintended bias. Companies, both developers and deployers, should use impact assessments as a tool for the responsible development and use of high-risk AI systems.

III. Distinguishing Between Different Entities in the AI Ecosystem

Much like privacy and security laws worldwide distinguish between different types of companies that handle consumers' personal data, AI laws should distinguish between different actors in the AI value chain, such as developers and deployers, to ensure that legal frameworks accurately assign obligations to a company based on its role in the AI ecosystem. A developer is the company that designs, codes, or produces an AI system, such as a software company that develops an AI system for speech recognition. A deployer, in contrast, is the company that uses an AI system, such as a bank that uses an AI system either developed internally or by a third party to make loan determinations. Each type of company will have access to different types of information about an AI system and will be positioned to take different actions to mitigate the risks associated with the AI system. AI policies that distinguish between these roles can ensure that the appropriate company within the various real-world AI supply chains can identify and mitigate risks.

Distinguishing between these two types of entities based on their role in the AI ecosystem can ensure companies are better able to fulfill their obligations and better protect consumers. For example, a developer would be able to describe the features of data used to train an AI system, but it generally would not have insight into how the AI system is used after another company has purchased and implemented the AI system. Instead, the deployer using the system is generally best positioned to understand how the AI system is being used, whether that use aligns with its intended use, whether and how to incorporate human oversight, the outputs from the AI system, any complaints received, and real-world factors affecting the system's performance.

We appreciate that SB 1047 recognizes the unique role developers play in the AI ecosystem, however, the bill does not account for AI deployers and holds developers of covered models responsible for obligations outside of their purview. For example, determining whether a covered model potentially has hazardous capabilities will largely depend on the context in which the covered model is deployed and mitigation measures undertaken by the deployer, which is often a different entity than the developer of the model. Requiring covered model developers to account for the specific harms included in the bill assumes developers have access to information they often do not and are not positioned to mitigate. Additionally, SB 1047 requires developers to incorporate full shutdown capabilities into their models and potentially holds developers liable for downstream uses over which they have no control. We recommend the bill be revised to account for AI developers and AI deployers and to place appropriate, role-based obligations on these different actors. We also recommend striking "or otherwise has responsibility for an artificial

[7] For example, three state privacy laws already require companies to conduct impact assessment for specific activities, including processing sensitive personal data, engaging in targeted advertising, or selling personal data; seven more state privacy laws will soon do so. Colorado, Connecticut, and Virginia already impose these requirements. See Colorado Privacy Act, Colo. Rev. Stat. Tit. 6, Art. 1, Pt. 13 §§ 6-1-1301–6-1-1313; Connecticut Data Privacy Act Conn. Gen. Stat. Tit. 42, Ch. 743jj, Sec. 42-515-525; Virginia Consumer Data Protection Act; Va. Code Tit. 59.1, Ch. 53, § 59.1-575-585. Recently passed state privacy laws in Florida, Indiana, Montana, Oregon, Tennessee, and Texas will also require impact assessments for certain activities. Globally, privacy and data protection laws worldwide use impact assessments as a tool for improving accountability.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

intelligence model" in the definition of developer to clearly focus on the companies creating or producing AI models.

IV. Clear and Consistent Enforcement and Compliance Mechanisms

Where legislation focuses on high-risk uses of AI and appropriately distinguishes between the roles and responsibilities of AI developers and AI deployers, we support exclusive enforcement by the Attorney General. Effective enforcement is important to protecting consumers, ensuring that businesses meet their obligations, and deterring potential violations.

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. Existing technical standards for AI are nascent and should be developed consistent with longstanding voluntary, market-driven, and consensus-based approaches to standards development. BSA supports the development of AI standards and best practices currently underway within international standards organizations and encourages California to adopt interoperable approaches to these AI standards and best practices as they mature.

Additionally, the bill requires developers of covered models to report AI safety incidents, which are vaguely defined, to the Frontier Model Division within 72 hours. Given the ambiguity surrounding the definition of AI safety incident, we are concerned this requirement would result in over-notification to the division and thereby not promote safety. The short timeline for reporting incidents also creates challenges for companies, which often will not have a comprehensive view of the scope of any incident within that limited time frame. Further, the notification requirement obligates companies to divert resources from addressing the safety of systems to fulfilling burdensome reporting requirements at a critical time. For these reasons, we suggest deleting this requirement.

* * *

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

CC: Members of the Senate Judiciary Committee; Senator Scott Wiener

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org