The Software Alliance

BSA

February 29, 2024

The Honorable James Maroney
The Honorable Michael D'Agostino
General Law Committee
Legislative Office Building, Room 3500
Hartford, CT 06106

Dear Chairs Maroney & D'Agostino:

BSA | The Software Alliance appreciates the opportunity to share insights on artificial intelligence (AI) and your work to promote AI innovation and protect against potential AI harms in SB 2. BSA is the leading advocate for the global software industry.[1] Our members are enterprise software and technology companies that create business-to-business products and services to help their customers innovate and grow. For example, BSA members provide cloud storage services, customer relationship management software, human resource management programs, identity management services, and collaboration software. Businesses entrust some of their most sensitive information with BSA members, and our companies work hard to keep that trust.

Enterprise software, including AI, is supporting digital transformation in every sector of the economy. AI is not just about robots, self-driving vehicles, or social media. It is used by businesses of all sizes to create the products and services they provide to consumers, to improve their internal operations, and to enhance their capacity to make data-informed decisions. BSA members are on the leading edge of providing businesses of all sizes and in every sector of the economy with the trusted tools they need to leverage the benefits of AI.[2]

As leaders in the development of enterprise AI, BSA members have unique insights into the technology's tremendous potential to spur digital transformation and the policies that can best support the responsible use of AI. BSA's views are informed by our recent experience working with member companies to develop 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

[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.

Framework outlines a lifecycle-based approach for performing impact assessments to identify risks of AI bias and highlights corresponding risk mitigation best practices.

We greatly appreciate your work in drafting SB 2 and the opportunity to provide feedback on the legislation. Our initial recommendations below focus on our core priorities in the introduced bill: (1) improving the provisions that address risks of high-risk AI systems, (2) identifying a range of practical concerns with the provisions on generative AI systems, and (3) recognizing the bill's approach to enforcement.

I. High-Risk AI Systems

We welcome the bill's focus on high-risk uses of AI systems. We believe it is important to prioritize AI policies that address activities that pose the most significant risks to consumers. The provisions in Section 2 and 3 create a strong foundation for addressing these risks. We have several recommendations about improving these provisions so that they work in practice.

a. **The consequential decision definition should be revised.** We appreciate that high-risk uses are tied to consequential decisions. To avoid overbroad application, we recommend focusing the definition of this term on eligibility determinations, changing "access to" to "eligibility for and results in the provision or denial of" in the definition.

b. **The list of information the developer provides to a deployer about a high-risk AI system should be revised to reflect the developer's role.** Section 2(b) outlines the information a developer must share with a deployer. However, it includes disclosure of an item that would not be within the purview of developers. Specifically, the bill requires developers to explain how a consumer can monitor the system once deployed. Because developers design, code, or produce AI systems, and deployers use AI systems, they have access to different types of information. In this instance, deployers are best positioned to provide information about how consequential decisions are made.

c. **The bill should require mitigation, not elimination, of known risks.** Section 3 of the bill requires implementation of a risk management program that, among other things, eliminates any known or reasonably foreseeable risks of algorithmic discrimination. While we support risk management programs, we note that elimination of all potential risks is not an attainable goal. Instead, risk management programs should enable the identification and mitigation of risks.

d. **The bill's duty of care approach should be reframed.** The bill requires developers and deployers to use reasonable care to avoid known or reasonably foreseeable risks of algorithmic discrimination, which can create amorphous duty-based obligations. In lieu of framing obligations under a duty of care, the bill should focus on steps that organizations can take to identify and mitigate risks. If the duty of care is retained, it should be narrowly tailored and described with more specificity.

e. **The bill's requirement for developers to report when an AI system has been used to cause algorithmic discrimination should be eliminated.** Section 2(e) requires developers to inform deployers and the Attorney General when they discover or are informed that a deployed high-risk AI system has caused algorithmic discrimination. As an initial matter, such a requirement envisions an ongoing postdeployment relationship with the deployer, which may not be the case. Further, one deployer's use of the high-risk system in a discriminatory manner does not render all other uses discriminatory, and such notice would often be irrelevant to another deployer's use of the system. We suggest striking this requirement.

f. **The bill should require developers to implement a risk management program.** We support the bill's requirement that deployers implement a risk management program. We also believe that the NIST AI Risk Management Framework is an important, flexible tool and support its reference as one way for implementing risk management programs. Both developers and deployers have responsibilities to address risks along the AI value chain. As a result, the bill's requirement to implement risk management programs should apply to developers too.

g. **We support the bill's flexibility in allowing the use of impact assessments conducted under other laws.** The bill provides that an impact assessment completed for another law or regulation will satisfy the bill's requirements if it is reasonably similar in scope and effect. This is an important provision, and we support its inclusion. Other laws, including some state privacy laws, require impact assessments, and authorizing the use of assessments that address similar issues will significantly ease compliance burdens.

h. **We support the bill's treatment of impact assessments.** We appreciate that the bill recognizes that impact assessments are rigorous internal processes, and that it requires disclosure to the Attorney General only in connection with an investigation. We also appreciate the bill's recognition that impact assessments contain confidential information and support the trade secret protection and Freedom of Information Act exemption.

II. Generative AI Systems

The bill's approach to regulating generative AI raises significant concerns, because it would impose a broad range of novel requirements that do not work in practice. As an initial matter, the bill's approach to generative AI is not risk-based and instead singles out a specific kind of technology to regulate, instead of focusing regulation on particular uses of the technology. Such an approach is overbroad and does not prioritize AI-related uses that pose the most significant risks to consumers.

The generative AI provisions would also impose sweeping restrictions across this technology. Unlike the EU AI Act, which requires providers of general purpose AI models to maintain technical documentation and includes additional obligations for large models that create systemic risks, the bill places a wide range of obligations on all general purpose AI systems, including a requirement to mitigate certain risks, data governance requirements, requirements relating to the appropriate level of performance, and obligations for authenticating, detecting, and labelling synthetic content. The bill also requires developers of generative AI systems to conduct impact assessments and protect consumers against a wide range of risks through a duty of care. These obligations go well beyond the treatment of general purpose AI models in the EU AI Act, which spurred significant debate among a range of stakeholders and policymakers.

Further, the bill requires a general purpose AI model's performance, interpretability, and safety to be assessed by independent experts. An inquiry into these aspects can implicate confidential or proprietary information, and companies should not be required to disclose this information to third parties. Moreover, AI auditing standards are nascent. In the absence of standards to audit against, there is a lack of objective benchmarks for auditors to use, which can lead companies to select auditors based on their own preferred criteria, undermining the goal of such audits. We believe that internal evaluation mechanisms are sufficiently robust to identify risks. We expect the bill’s provisions on generative AI to generate significant and widespread concerns from industry stakeholders, which risks slowing down the important parts of the bill focused on high-risk AI systems.

III. Enforcement

Regarding the sections on high-risk and generative AI systems, exclusive enforcement by the Attorney General ensures a consistent approach to enforcement. We appreciate that these portions of the bill do not create a private right of action and expressly state they do not create a private right of action under any other law.

* * *

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 SB 2 as it progresses through the legislature.

Sincerely,

Meghan Pensyl

Meghan Pensyl
Director, Policy