May 15, 2023

The Honorable Chris R. Holden
Capitol Office
1021 O Street
Suite 8220
Sacramento, CA 94249

Dear Chair Holden:

BSA | The Software Alliance appreciates the opportunity to share insights on Assembly Bill 331, which seeks to address concerns about bias in AI systems. This is an important issue for BSA | The Software Alliance and our member companies, and we support the goal of preventing unlawful discrimination in consequential decision systems.

BSA is the leading advocate for the global software industry.[1] Our members are enterprise software companies that create business-to-business technologies that help other businesses innovate and grow.[2] For example, BSA members provide tools including cloud storage services, customer relationship management software, human resource management programs, identity management services, and collaboration software. BSA members are on the leading edge of providing AI-enabled products and services. As a result, they 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 Framework outlines a lifecycle-based approach for performing impact assessments to identify risks of AI bias and highlights corresponding risk mitigation best practices.

The approach taken in AB 331 aligns well with the BSA Framework and includes key elements that we support, including those highlighted below. We also have suggestions for how the bill can be improved, and we welcome the opportunity to work with you as you consider changes to the bill.

---

[1]BSA's members include: Adobe, Alteryx, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, Databricks, DocuSign, Dropbox, Graphisoft, IBM, Informatica, Juniper Networks, Kyndryl, MathWorks, Microsoft, Okta, Oracle, Prokon, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Splunk, Trend Micro, Trimble Solutions Corporation, TriNet, Twilio, Unity Technologies, Inc., 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

Our comments focus on five aspects of the legislation.

First, BSA supports the legislation's recognition of the different roles and responsibilities of developers and deployers of AI systems.

Just as privacy and security laws distinguish between different types of companies that handle consumers' personal information, distinguishing between developers and deployers[4] ensures that legal frameworks accurately assign obligations to a company based on its role in the AI ecosystem. As a result, companies are better able to fulfill those obligations and better protect consumers.

For example, the developer of an AI system is generally well-positioned to describe the operation of that AI system, but it would not typically have insight into how the AI system is used after another company has purchased and implemented the AI system. In contrast, the deployer using an AI system is generally best positioned to understand how the AI system is being used, to understand whether that use aligns with the intended uses of that AI system, to address whether and how to incorporate human oversight of the AI system, to assess outputs from the AI system, to address any complaints received, and to understand real-world factors affecting the system's performance.

BSA supports AB 331's inclusion of this important distinction by separately defining "developers" and "deployers" and by creating obligations for both types of companies that reflect their different roles.

Second, BSA supports the legislation's requirement for companies that develop and use AI systems for consequential decisions to conduct impact assessments and design evaluations.

BSA supports the overarching goal of AB 331, which is to ensure high-risk uses of AI are subject to safeguards. One crucial safeguard that promotes responsible uses of AI systems is ensuring that companies that develop or use high-risk AI systems establish a comprehensive approach for performing impact assessments and design evaluations. Impact assessments are widely used in a range of other fields — from environmental protection to data protection — as an accountability mechanism that promotes trust by demonstrating that a system has been designed in a manner that accounts for the potential risks it may pose.

BSA supports requiring companies to conduct impact assessments and design evaluations for AI systems used to make consequential decisions. These assessments and evaluations are important accountability tools that help businesses identify, document, and mitigate AI risks. Notably, they are also helpful tools in detecting and mitigating potential bias that could result in unlawful discrimination. Any legislation creating impact assessments and design evaluations should apply to high-risk uses and clearly distinguish requirements for developers and deployers.

BSA supports AB 331's approach of creating separate obligations for developers and deployers to conduct impact assessments and design evaluations for automated decision tools that make consequential decisions.

- *Timing for impact assessments*. In addition to annual impact assessments, the bill requires companies to conduct a new assessment every time there is "any significant

[4] *See* BSA, AI Developers and Deployers: An Important Distinction, *available at* https://www.bsa.org/files/policy-filings/03162023aidevdep.pdf.

update." In practice, this could be read to require such frequent updates it would reduce incentives for companies to conduct thorough assessments, undermining the bill's objective. Instead, we recommend requiring a new assessment if there are material changes to the purpose for which an automated decision tool is used. This creates a clear trigger for additional assessments and ensures new impact assessments are conducted if a tool will be used for a new purpose, without requiring new assessments for updates that merely improve functionality.

- *Focus of Developer's Impact Assessment.* The bill requires a developer's impact assessment to include "a summary of the type of data collected from natural persons and processed by the automated decision tool." However, this requirement refers to information possessed by the deployer of the system — and that information is often unavailable to the company that developed a system. Instead, the developer's obligation should focus on providing an overview of the type of data it used to train the automated decision tool, rather than data that a deployer will collect during the tool's later use. Ensuring that these obligations are tailored to each entity's role will help the bill's safeguards function in practice.

- *Adverse Impact Analysis*. The bill's requirement to conduct an adverse impact analysis presumes that companies have or should have access to data needed for such an analysis (e.g., a bank having information on a customer's genetic status, which would be needed to test a tool for genetic discrimination in credit decisions). We recommend that this provision be revised to require an "assessment for the reasonably foreseeable risks of algorithmic discrimination" and to clarify that the assessment be appropriate to the data to which a developer or deployer has access.[5]

Third, BSA recommends that the definition of "consequential decision" be more narrowly tailored to provide clear guidance of what conduct is covered under the bill and focuses on activities that pose a high risk to individuals.

BSA supports linking obligations to consequential decisions, as AB 331 does, but that term should be defined in a way that gives companies sufficient notice of the types of decisions governed by the law. Currently, the bill defines the term as a "decision or judgment that has a legal, material, or similarly significant effect on an individual's life relating to the impact of, access to, or the cost, terms, or availability of" an extensive list of enumerated categories. We recommend defining consequential decision to focus more narrowly on determinations that have the highest risk to individuals and meet a greater threshold than "relating to the impact of" particular areas. We agree with an approach that focuses on legal or similarly significant effects, which should be defined as decisions that determine "eligibility for" or result in the provision or denial of important services, e.g., housing, employment, education, healthcare, physical places of public accommodation, and insurance.

Although AB 331 identifies similar categories, among others, including the phrase "relating to the impact of" may result in overbroad application in practice. For example, the current language could sweep in automated tools that merely help with appointment scheduling for healthcare providers because that function is "relating to" the "availability" of a healthcare service. A more nuanced approach would appropriately focus on instances in which a

[5] The standard that should be applied should also align with guidance provided by the Equal Employment Opportunity Commission. *See, e.g.*, EEOC, The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees, *available at* https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence.

provider's use of an automated decision tool results in the provision or denial of care in a particular scenario.

Moreover, AB 331 defines the categories themselves too broadly. For example, it is reasonable to conclude that consequential decisions in employment include hiring decisions, but the proposed bill would also include task allocation, which can sweep far more broadly without raising the same level of risk to individuals. An automated decision tool used to schedule shifts for fast food workers should not be subject to the same requirements as the use of a tool that is a controlling factor in making a hiring decision. Further, some categories identified in the bill, such as "[a]ccess to benefits or services," are undefined and could have unintended consequences. Under the current language, a decision to use AI in advertising for a rewards program could be construed as "access to a benefit or service," yet it does not pose a high risk to individuals. The same unintended consequence applies to the cybersecurity space, where AI can be used to detect cyber threats and can help enhance the cybersecurity posture of an organization by, for example, granting or denying an individual's access to his financial aid benefits based on the risk level of an IP address he is logging in from. In this case, "access" from a cybersecurity standpoint is wholly different from "access" from an eligibility standpoint.

In sum, the Assembly should narrow the definition of consequential decision to clearly identify high-risk use cases included within its scope. Specifically, we recommend amending the definition of consequential decision by replacing the phrases "impact of," "access to," and "availability of" with more specific language such as "eligibility for."

Fourth, BSA supports the legislation's requirement for developers and deployers to implement a governance program.

A governance program provides the overarching framework necessary to identify, document, and mitigate AI risks. It ensures that appropriate personnel have been designated to oversee accountability measures, that organizational policies are established to guard against risks of algorithmic discrimination, and that processes are in place to implement safeguards that address any issues identified in the impact assessments and design evaluations. We support the bill's recognition of the important role of these functions. We also support the bill's reference to mapping, measuring, managing, and governing risks, which highlights the functions articulated in the National Institute of Standards and Technology's AI Risk Management Framework (RMF). The AI RMF is an important accountability tool and can serve as a useful guide for organizations aiming to address AI risks.

With respect to the specific program requirements, we recommend that in lieu of a two-year retention requirement for impact assessments, the bill should instead direct companies to preserve them for a reasonable period of time in light of the intended use. This would allow more flexibility to tailor retention activities to the particular circumstances.

BSA recommends that any legislation requiring impact assessments ensure that those requirements are enforced on a timeline that provides businesses time to create strong governance programs. In some cases, a company may act as both a developer and a deployer and will therefore need to develop two distinct compliance plans. It is critical that these programs are developed with ample time to construct a thorough governance program, to effectuate the goals of AB 331. We strongly encourage providing companies with two years between the time a bill is signed into law and its effective date. We therefore encourage you to extend the effective date past January 1, 2025, to allow time for more effective compliance.

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

Finally, BSA recommends strong and exclusive regulatory enforcement.

Strong enforcement is needed in any legislation that requires companies to develop and use high-risk AI systems in trustworthy ways. In our view, AB 331 should be exclusively enforced by a strong statewide regulator that can establish clear guidance and a consistent approach to enforcing the bill's requirements. Exclusive governmental enforcement by a single regulator ensures companies know how to implement AB 331's obligations — and avoids the conflicting interpretations and confusion likely to arise if courts reach different conclusions about how companies are to apply the bill's obligations.

BSA appreciates the recent amendments to 22756.8 to limit civil actions against a deployer or developer for violations of the bill by designating that power exclusively with public attorneys. We note that this provision could be further improved by consolidating the disparate governmental enforcement efforts within a single governmental entity, e.g., the Attorney General's office. We believe this change will further increase consistency in enforcement.

We also appreciate the amendments adding a right to cure violations before the public attorneys can file suit. We understand the assurances provided by requiring a statement that the violation has been cured, but that should not extend to potential future violations. It would be impossible to guarantee that no issues will arise in the future and unreasonable to impose liability for perjury if a subsequent violation occurs.

We also remain concerned with the private right of action established under 22756.6. Under that provision, an individual may bring civil actions against a deployer if a deployer uses an automated decision tool that *results in* algorithmic discrimination. We agree with the overarching goal of this provision: that AI tools should not be used to unlawfully discriminate against individuals. Indeed, at the federal level we have repeatedly called on the Administration to ensure that anti-discrimination laws remain fit for purpose in the digital age.[6] However, enforcing AB 331's requirements through a private right of action would not just create the potential for individuals to file significant amounts of lawsuits against consumer-facing companies to litigate the new standard created by AB 331. It also encourages companies to resolve disputes about these unclear obligations through rounds of secondary business-on-business litigation, shifting those unclear obligations from one company to another. That doesn't help consumers, who would be better served by consistent enforcement of this new standard by an agency that issues clear guidance that companies can readily implement. We therefore believe a strong, centralized approach to enforcement is the best way to develop sound practices. For this reason, we strongly recommend that the private right of action included in the current version of the legislation be removed.

In addition, we recommend AB 331 avoid requiring companies to proactively disclose impact assessments and design evaluations to a regulator, but instead recognize that a regulator may appropriately request those materials through its existing authorities. This is the approach taken in many state-level and international privacy laws, where companies are required to assess certain data processing activities but are not required to proactively provide those assessments to a regulator. Taking the same approach here will help ensure companies are incentivized to conduct robust assessments while ensuring a regulator can request these materials in its enforcement role.

[6] *See*, e.g., BSA Comments to NTIA on Privacy, Equity, Civil Rights (March 6, 2023), *available at* https://www.bsa.org/files/policy-filings/03062023ntiapriveq.pdf; BSA, Submission Regarding OSTP AI Bill of Rights Initiative (Jan. 13, 2022), *available at* https://www.bsa.org/files/policy-filings/01132022ostpai.pdf.

Consumers are best served by a clear enforcement of statutory obligations that can be readily implemented. We therefore support a strong, centralized approach to enforcing AB 331 with exclusive agency enforcement. This approach can implement the legislation in a manner that encourages companies to invest in engineering that protects Californians in line with clear regulatory actions and guidance.

* * *

Once again, thank you for the opportunity to provide the enterprise software sector's perspective on AB 331. BSA is focused on supporting safeguards on the use of AI that enhance trust in these technologies. We look forward to working with you on the proposed legislation in further detail and serve as a resource as you continue to consider these important issues.

Sincerely,

Matthew Lenz
Senior Director and Head of State Advocacy
BSA | The Software Alliance

CC: Assemblymember Bauer-Kahan, Sponsor

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org