June 13, 2024

The Honorable Tom Umberg
1021 O Street
Suite 6530
Sacramento, CA 942849-0070

Dear Chair Umberg:

BSA | The Software Alliance appreciates the opportunity to share insights from the enterprise software sector on artificial intelligence (AI) and AB 2930. 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. 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 uses of AI. BSA's views are informed by our experience with members developing 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.

While the approach taken in AB 2930 aligns well with the BSA Framework and includes key elements that we support, recent amendments to the bill create new concerns. At the outset, we appreciate that some recent amendments help ensure the bill's workability in practice, like

[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

exempting cybersecurity-related technologies, which helps ensure organizations have flexibility in identifying and mitigating cybersecurity risks. However, we're concerned other provisions detract from the bill's workability and conflate the roles of different actors in the AI value chain. We welcome the opportunity to work with you as you consider further changes to the bill.

Our comments focus on five aspects of the legislation:

- Distinguishing among different actors in the AI value chain;
- Conducting impact assessments;
- Narrowing the scope of the definition of consequential decision;
- Implementing a governance program; and
- Ensuring exclusive regulatory enforcement.

**I.** The bill should be revised to better reflect the different roles and responsibilities of different actors in the AI value chain.

BSA supports the bill's recognition of the different roles of different entities in the AI value chain. Just as privacy and security laws distinguish between different types of companies that handle consumers' personal information, distinguishing between AI developers and AI 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.[5] BSA supports AB 2930'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.

The bill should be revised in two ways to better reflect the distinction between the roles of developers and deployers.

- First, a new provision introduced in the Assembly Privacy and Consumer Protection Committee that prohibits developers from providing to deployers an automated-decision tool that results in algorithmic discrimination blurs the lines between these different actors. Both developers and deployers should have obligations to protect against algorithmic discrimination, but those obligations must be based on each entity's distinct role to be workable in practice. As written, this provision may impose liability on a *developer* when a *deployer* chooses to use an automated decision tool in a discriminatory way–even though the developer does not control how the deployer uses the tool. We strongly recommend removing this provision, to ensure the bill creates role-based obligations that assign each type of company obligations that reflect the information that company has and the action that company can take to identify and mitigate potential harms.

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

[5] 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.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

- Second, where the bill requires developers to provide deployers with a statement regarding the intended uses of the automated decision tool and certain documentation about the automated decision tool, we appreciate that developers are not required to disclose trade secrets. We recommend this exception be expanded to include both intellectual property and confidential information, to better protect developers' proprietary information.

**II.** BSA supports requiring companies that develop and use AI systems for consequential decisions to conduct impact assessments.

BSA supports the overarching goal of AB 2930, 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 deploy AI systems for high-risk uses 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 AB 2930's general approach of creating separate obligations for developers and deployers to conduct impact assessments for automated decision tools that make consequential decisions.

We recommend addressing several aspects of the bill's impact assessment obligations:

- *Broad New Exceptions Undermine Bill's Safeguards.* Recent amendments create a broad set of exceptions to the bill's impact assessment obligation for deployers. We are concerned that these exceptions weaken the bill's protections and may inadvertently cause companies to deploy high-risk AI systems without appropriately assessing the potential harms associated with their particular use. We recommend that these exceptions at least are narrowed to make clear that deployers may have existing legal obligations under other laws to protect against discrimination.

- *Frequency of Impact Assessments*. In addition to annual impact assessments, the bill requires companies to conduct a new assessment every time there is "any significant update," meaning "a new version, new release, or other update . . . that materially changes its principal use, principal intended use, or outcome." We appreciate that this term has been narrowed to focus on material changes to the purpose for which an automated decision tool is used. We recommend maintaining this language, which 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, that information is often unavailable to the company that developed a system—and would instead generally be available to the deployer using the system. Instead of this requirement, 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 focusing on 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.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

- *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.[6]

- *Promoting Interoperability.* As AI rapidly evolves and is integrated into our daily lives and business processes, companies will conduct impact assessments for uses in California and other jurisdictions. Interoperability is crucial to ensuring that best practices and norms can be leveraged across geographies. We recommend that AB 2930 incorporate an interoperable approach to impact assessments. Specifically, the legislation should include a requirement that "[i]mpact assessments conducted by a deployer or developer for the purpose of compliance with other laws, regulations, or generally accepted industry framework, such as those developed by the National Institute of Standards and Technology, may comply [under this section] if the assessments have a reasonably comparable scope and effect."

**III.** BSA recommends that the definition of "consequential decision" be narrowed to provide clear guidance of what conduct is covered under the bill and to focus on activities that pose a high risk to individuals.

BSA supports linking obligations to consequential decisions, as AB 2930 does. However, that term should be defined in a way that gives companies clear 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 access to government benefits or services, assignments of penalties by the government, or the impact of, 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 and result in the provision or denial of" important services, e.g., housing, employment, education, healthcare, physical places of public accommodation, and insurance.

We recommend narrowing AB 2930's definition of consequential decision in three ways:

- First, the phrase relating to "the impact of, or the cost, terms, or availability of" should be revised to avoid overbroad application of this definition. 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. The bill should adopt a more nuanced approach, by focusing on instances in which a provider's use of an automated decision tool results in the provision or denial of care in a particular scenario.

- Second, AB 2930 defines the categories themselves too broadly. For example, in

[6] 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.

addressing employment-related use of AI, the bill includes certain instances of task allocation, which can sweep far more broadly than may be intended. 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 substantial factor in rejecting an applicant from the hiring process. As a result, we recommend deleting task allocation from the list of examples in the definition of consequential decision. If task allocation is retained, then we recommend revising it by adding "automatically" before "limits" and adding "based on individual performance or behavior" after "employees" to narrow the scope.

- Third, "access to" government benefits or services is overly broad and could have unintended consequences. For example, a company that provides a customer data platform may contract with a local government agency to promote a new government benefit program. The government agency may use the company's AI tools to create and send notifications to individuals encouraging them to sign up for the new benefit. Even though the government agency created this notification and decided which individuals should receive it, the bill may hold the customer data platform company accountable for affecting "access to" the government benefit.

In sum, the definition of consequential decision should be narrowed to clearly identify high-risk use cases included within its scope. Specifically, we recommend:

- (1) amending the definition of consequential decision by replacing the phrases "impact of" and "the costs, terms, or availability of" with more specific language such as "eligibility for and results in the provisions or denial of" and

- (2) narrowing the enumerated categories, including by deleting task allocation from the list of examples.

**IV.** 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 five-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 also 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 2930. We appreciate the stagged compliance deadlines established for state government deployers and strongly encourage providing companies with at least 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, 2026, to allow time for more effective compliance.

**V.** 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 2930 should be exclusively enforced by the Attorney General, who 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 2930'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.

The legislation allows public attorneys and the Civil Rights Department, in addition to the Attorney General, to bring civil actions against a deployer or developer for violations of the bill. We note that this provision could be improved by consolidating the disparate governmental enforcement efforts within the Attorney General's office. We believe this change will further increase consistency in enforcement. Further, we recommend that courts should not be able to award reasonable attorney's fees and litigation costs in civil actions, as such an option would be unduly punitive to defendants.

We appreciate that companies have the right to cure violations before the public attorneys or the Civil Rights Department can file suit. We understand the assurances provided by requiring a statement that the violation has been cured, and developers and deployers should certainly take steps to ensure the violation is cured, but it is unreasonable to impose liability for perjury. Several unforeseen events could arise, and imposing a criminal penalty is disproportionate to the violation.

Additionally, while we understand that the Civil Rights Department should be able to investigate reports of algorithmic discrimination and request companies' impact assessments in connection with those investigations, we recommend that companies be required to provide impact assessments within 30 days of the Civil Rights Department's request, not seven days. This change allows sufficient time to redact proprietary information, like trade secrets, which this section of the legislation helpfully protects from public disclosure. We recommend the bill's safeguards surrounding trade secrets be expanded to include intellectual property and confidential information to better protect proprietary information.

* * *

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

Sincerely,

Meghan Pensyl

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

Meghan Pensyl
Director, Policy

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org