August 11, 2025

The Honorable Anna Caballero
Chair, Senate Appropriations Committee
California State Capitol, Room 412
Sacramento, California 95814

Re: Assembly Bill 1018 (Bauer-Kahan) - Automated Decision Systems - OPPOSE

Dear Chair Caballero, Vice Chair Seyarto, and Members of the Senate Appropriations Committee,

The Business Software Alliance appreciates the opportunity to share insights from the enterprise software sector on artificial intelligence (AI) and Assembly Bill 1018, 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.

As you consider how to regulate AI through AB 1018, we want to underscore the importance of ensuring any AI legislation creates thoughtful, clear guardrails for companies and protects consumers. To achieve this, we strongly recommend AI legislation:

- Focus on the uses of AI that have the greatest impact on consumers;
- Reflect the different roles and responsibilities of different actors along the AI value chain;
- Require risk management programs and impact assessments;

[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

- Ensure strong enforcement; and
- Promote interoperability and incorporate stakeholder feedback.

While we support ensuring that companies that develop and use AI systems to make consequential decisions about consumers have obligations to develop and use those AI systems responsibly, we are concerned that AB 1018 is unworkable in practice and oppose the legislation unless it is amended to improve its workability.

Below we offer feedback on several aspects of AB 1018 and AI legislation generally, and we have views on other provisions of AB 1018 as well. We would welcome the opportunity to discuss our feedback with you or a member of your staff.

I. AI Creates Benefits Across Industry Sectors

As you consider how to regulate AI, it is important to recognize that the economic benefits of AI are not limited to one industry sector or one business model. Instead, the promise that AI may one day impact every industry is quickly turning into a commercial reality and driving digital transformation across sectors. Airlines now use AI systems to more efficiently clean planes between flights; farmers use AI to analyze large amounts of weather information to maximize their harvest; manufacturers use AI to test new prototypes; and construction companies build AI-generated "digital twins" of real-life cities to understand the impacts of a proposed design.

Companies in all industries use AI-powered enterprise software, including:

- In healthcare, a large pharmacy chain uses an advanced platform to forecast demand and redistribute medications across thousands of store locations and to deliver near real-time insights and recommendations for pharmacists to provide more personalized advice to patients. This helps managers understand the supply chain, store labor and productivity, patient vaccine scheduling, and prescription pickup processes.

- In manufacturing, a car maker used generative AI technology to redesign a seat bracket, which secures seat belt fasteners to seats and seats to floors, that is 40 percent lighter and 20 percent stronger than the previous iteration. Changes like these can help reduce the amount of material needed to build a car and make vehicles more fuel efficient.

- In agriculture, the research division of an enterprise software provider partnered with a climate risk company to develop software capable of providing more accurate long-range weather predictions. Traditional weather forecasting methods can provide accurate predictions for a seven-day window. By leveraging AI, the

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

researchers are developing new forecasting models to provide accurate predictions of weather trends two to six weeks out from a given date. By providing reliable extended forecasts, these tools will help water managers predict snowpack and water availability for irrigation, hydropower, and other critical agricultural and environmental uses.

Because BSA members work with companies across every sector of the economy, we have unique insights into AI's tremendous potential to further spur digital transformation and the policies that can best support the responsible use of AI. BSA's views are informed by our 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 and highlights corresponding risk mitigation best practices.

II. AI Legislation Should Focus on High-Risk Uses of AI

AI legislation should focus on the uses of AI that have the most impact on consumers' lives. 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.

However, when AI systems are used to decide whether consumers are granted or denied important benefits and services, like housing, healthcare, and employment opportunities, companies should be accountable for developing and deploying those systems responsibly.

We are concerned the scope of AB 1018 is so broad that it extends beyond high-risk AI systems to even mundane uses of software. Below we suggest improvements to several of the bill's definitions to help tailor the legislation to the uses of AI that have the most significant impacts on consumers.

We recommend AB 1018 use the following terms:

- *High-risk artificial intelligence system:* an artificial intelligence system designed or substantially modified with the intended purpose of making consequential decisions about an individual.

[3] *See* BSA's Framework to Build Trust in AI, *available at* https://ai.bsa.org/confronting-bias-bsas-framework-to-build-trust-in-ai.

- *Artificial intelligence system:* A machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different artificial intelligence systems vary in their levels of autonomy and adaptiveness after deployment.

- *Consequential decision*: a decision made by a deployer that has a material, legal, or similarly significant effect on a consumer's eligibility for, and results in the provision or denial of, housing, employment, credit, education, access to physical places of public accommodation, healthcare, or insurance.

- *Intentionally and substantially modified or intentional and substantial modification* means the following:
    - (1) for a non-high-risk AI system, an entity intentionally and substantially modifies that system when the entity intentionally changes or establishes its intended purpose so that it becomes a high-risk AI system, whether through development or deployment; or
    - (2) for a high-risk AI system, an entity intentionally and substantially modifies that system when the entity intentionally changes the core capabilities of the high-risk AI system in a way that materially increases the reasonably foreseeable risks of algorithmic discrimination in connection with that system's use in making a consequential decision.

III. AI Legislation Should Require Risk Management Programs and Impact Assessments, Not Third-Party Audits

For high-risk uses of AI, new safeguards are important—and legislation can leverage tools that already exist to help companies identify and mitigate potential risks. BSA supports requiring companies that develop or deploy AI for high-risk uses to: (1) adopt risk management programs; and (2) conduct impact assessments. These measures can help companies identify and mitigate risks when AI makes important decisions about consumers—and increase trust that AI is developed and used responsibly. In particular:

- 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. One way for companies to establish risk management programs is by using the AI Risk Management Framework (AI RMF), developed by the National Institute of Standards and Technology (NIST). The AI RMF builds on NIST's work creating frameworks for managing cybersecurity and privacy risks. Ultimately, effective AI risk management

programs should support cross-company coordination to promote the identification and mitigation of risks across the lifecycle of an AI system.

- Impact assessments are a key part of a meaningful risk management program. Both developers and deployers should use impact assessments as a tool for the responsible development and use of high-risk AI systems—and each type of company should conduct an impact assessment that reflects their role in developing or deploying the AI system. 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 that accounts for the potential risks it may pose to the public. Because impact assessments already exist today in other fields, they can be readily adapted to help companies identify and mitigate AI-related risks.[4]

We strongly recommend that AB 1018 be amended to focus on risk management programs and impact assessments and remove requirements for third-party audits. AB 1018 would require developers of automated decision systems to contract with third-party auditors to assess the developer's compliance with the legislation, consider and attempt to incorporate feedback from such auditors, and respond to auditors contracted by a deployer regarding the performance of the automated decision system. Such requirements would establish a sweeping audit regime—which creates significant concerns because the AI audit ecosystem is still immature and not fully equipped to assure the performance of AI systems.[5] Unlike regulated industries like accounting and financial services, there is not a professional body to certify AI auditors. As a result, the range of qualified and available auditors may vary greatly—and undermine the bill's goal of implementing audits consistently across different AI systems.

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

IV. AI Legislation Should Distinguish Between Different Entities in the AI Ecosystem

[4] *See* BSA, Impact Assessments: A Key Part of AI Accountability, *available at* https://www.bsa.org/policy-filings/impact-assessments-a-key-part-of-ai-accountability.
[5] *See* BSA, Enhancing AI Accountability: Effective Policies for Assessing Responsible AI, *available at:* https://www.bsa.org/policy-filings/enhancing-ai-accountability-effective-policies-for-assessing-responsible-ai.

The AI supply chain is evolving, and AI legislation should not create one-size-fits-all requirements when companies have very different roles.

All companies that develop and use high-risk systems have responsibilities to manage AI risks, but those obligations must reflect the role of each type of company, since each will know different information about an AI system and will be able to take different actions to identify and mitigate risks. Legislation should reflect these differences, to create obligations that work in practice to safeguard consumers.

While we appreciate that AB 1018 differentiates between developers and deployers of automated decision tools, we are concerned the bill distorts the respective responsibilities of each type of company and imposes unworkable obligations on developers. Below we suggest improvements to the definitions of developer and deployer and discuss our concerns with the bill's substantive obligations.

We recommend adopting the following definition in legislation on high-risk AI:

- *Developer of a high-risk AI system*: An entity doing business in the state that makes an artificial intelligence system publicly available for use in the state and that:
    - designs an AI system specifically intended to be a high-risk AI system;
    - intentionally and substantially modifies a high-risk AI system; or
    - intentionally and substantially modifies a non-high risk AI system so that it becomes a high-risk AI system.

- *Deployer*: An entity doing business in the Commonwealth that uses a high-risk AI system to make a consequential decision.

AB 1018 should also clarify that: (1) When an entity intentionally and substantially modifies a high-risk artificial intelligence system and is therefore the developer of that system, the entity is responsible only for actions related to the risk the entity introduces and information within its possession, custody, or control; and (2) When an entity acts as the deployer of a high-risk artificial intelligence system, the entity is responsible only for actions related to the risk the entity introduces and information within its possession, custody, or control.

Additionally, we are concerned the obligations imposed on developers are unworkable, because AB 1018 requires developers to account for risks they do not introduce and for information they often do not have access to. We strongly recommend that the bill's obligations on developers be significantly revised. We identify below several initial suggestions that can improve how the legislation works in practice.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

First, the bill should avoid distorting the roles of developers and deployers. For example, the bill requires developers to identify "developer-approved uses," which are defined to include information about the "deployment context" for AI systems they develop.[6] But it is the *deployer* of an AI system—not the developer—that decides the context for using an AI system in practice. Requiring developers to know the context in which deployers use their AI systems distorts the different roles of these different companies along the AI value chain. More broadly, we recommend the bill focus on evaluating actual uses of the automated decision system by the deployer, rather than relying on intended uses by the developer. We strongly recommend amending the definition of "developer-approved uses" to avoid distorting these roles.

Second, the obligations on developers should be revised to avoid sweeping and overly prescriptive requirements. AB 1018 would require developers, in performance evaluations, to assess the expected performance of the automated decision system for each developer-approved use, including the expected accuracy and reliability and any reasonably foreseeable effects of fine tuning, whether disparate treatment is intended to occur, and whether any disparate treatment is likely to occur, among other requirements. But the expected accuracy of an AI system depends in large part on how it is deployed—which is information known by the deployer, not the developer. The bill further distorts the roles of developers and deployers by requiring developers to consider potential disparate impacts and disparate treatment. But developers will likely have little to no insights into those issues, since developers often have no direct relationships with consumers whose information is processed by the AI system after it is deployed. We strongly recommend AB 1018's requirements on developers be amended to focus on: (1) impact assessments, rather than performance evaluations, as discussed above; and (2) on unlawful algorithmic discrimination, rather than disparate impacts or disparate treatment, to ensure that the bill creates clear guardrails for consumers and clear obligations for companies.

V. AI Legislation Should Ensure Strong Enforcement

We strongly recommend AB 1018 be amended to grant the Attorney General exclusive enforcement authority.

Strong enforcement is needed in any AI legislation. Amending AB 1018 to give the Attorney General exclusive enforcement authority can help that office establish clear guidance and a consistent approach to enforcing the bill's requirements. Currently, AB 1018 would be enforced by consumers through private rights of action under existing state laws and through various state and local authorities, namely:

[6] The bill defines developer-approved use as "a deployment context in which a developer intends a covered [automated decision system] to make or facilitate consequential decisions," including any reasonably foreseeable fine tuning of the covered automated decision system.

- The Attorney General;
- Certain district attorneys, county counsels, and city attorneys;
- Certain city prosecutors;
- The Civil Rights Department; and
- The Labor Commissioner.

Exclusive governmental enforcement by a single regulator ensures companies know how to implement the legislation's obligations—and avoids the conflicting interpretations and confusion likely to arise if courts reach different conclusions about how companies are to apply AB 1018's obligations.

VI. AI Legislation Should Promote Interoperability and Incorporate Stakeholder Feedback

California is home to global companies, and your legislation will be most effective when it is interoperable with other approaches to AI regulation. Global companies can better serve their customers when they build strong compliance programs that work across markets. We also encourage you to continue working with stakeholders as you develop your legislation, to understand how your AI law will work in practice, across a range of different industries and uses.

* * *

Thank you for allowing us to provide the enterprise software sector's perspective on AB 1018. We welcome the opportunity to further engage with you or a member of your staff on these important issues.

Sincerely,

Meghan Pensyl
Director, Policy