



**BSA Comments on OMB Memorandum on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence  
December 5, 2023**

BSA | The Software Alliance appreciates the opportunity to provide comments on the Office of Management and Budget's Memorandum on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (OMB memo).

BSA is the leading advocate for the global software industry.<sup>1</sup> BSA members are at the forefront of developing cutting-edge services — including AI — and their products are used by businesses across every sector of the economy.<sup>2</sup> For example, BSA members provide tools including cloud storage and data processing 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. BSA members have also supported the federal government's IT modernization efforts and have extensive experience providing software services, including AI, to the federal government. 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 experience working with member companies to develop the BSA Framework to Build Trust in AI,<sup>3</sup> a risk management framework we published more than two years ago to help companies mitigate the potential for unintended bias in AI systems. 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 best practices.<sup>4</sup> Our experience on these issues informs our recommendations on the OMB memo.

---

<sup>1</sup> BSA's members include: Adobe, Alteryx, Asana, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, Databricks, DocuSign, Dropbox, Elastic, Graphisoft, Hubspot, IBM, Informatica, Juniper Networks, Kyndryl, MathWorks, Microsoft, Okta, Oracle, Palo Alto Networks, Prokon, 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.

<sup>2</sup> See BSA | The Software Alliance, *Artificial Intelligence in Every Sector*, available at <https://www.bsa.org/files/policy-filings/06132022bsaaieverysector.pdf>.

<sup>3</sup> 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>.

<sup>4</sup> BSA has testified before the United States Congress and the European Parliament on the Framework and its approach to mitigating AI-related risks. See, e.g., Testimony of Victoria Espinel, Public Hearing on AI & Bias, Special Committee on Artificial Intelligence in a Digital Age, European Parliament, Nov. 30, 2021, available at [https://www.europarl.europa.eu/cmsdata/244265/AIDA\\_Verbatim\\_30\\_November\\_2021\\_EN.pdf](https://www.europarl.europa.eu/cmsdata/244265/AIDA_Verbatim_30_November_2021_EN.pdf); Testimony of Victoria Espinel, The Need for Transparency in Artificial Intelligence, Before the Senate Committee on Commerce, Science, and Transportation Subcommittee on Consumer Protection, Product Safety, and Data Security, available at <https://www.bsa.org/files/policy-filings/09122023aitestimonyoral.pdf>.

The OMB memo has important ramifications for federal IT modernization. Given the increase of AI in software services, it is critical that the guidance strikes the right balance in creating safeguards that promote the responsible use of AI while maximizing its benefits to the federal government. If it fails to do so, the gap between the commercial and federal marketplace will only grow.

The OMB memo establishes a range of important risk management practices across federal agencies. We are encouraged that the OMB memo requires agencies to take important steps to adopt strong risk management practices and attempt to categorize risks associated with government use of AI. This approach aims to help agencies focus on uses of AI that pose high risks to the public.

We recommend revising several aspects of the OMB memo, to ensure agencies can better implement these risk management practices. Specifically, we recommend that OMB:

- Continue to remove barriers to the responsible use and adoption of AI;
- Ensure uniform definitions of rights-impacting and safety-impacting AI across federal agencies;
- Ensure the definitions and list of rights-impacting and safety-impacting AI use cases focus on those risks that impact individuals;
- Further encourage the use of the NIST AI Risk Management Framework;
- Refine the minimum risk management practices required for federal agencies;
- Extend the timeline for agency implementation;
- Revise procurement obligations, to ensure vendors' ability to further train AI models and to recognize the importance of internal testing; and
- Harmonize concurrent regulatory actions affecting AI procurement.

**I. The OMB Memo Should Continue to Strengthen AI Governance and Remove Barriers to the Responsible Use of AI by Agencies.**

At the outset, we want to recognize that the OMB memo includes important guidance for strengthening federal agencies' AI governance and advancing AI innovation across the federal government. We appreciate the OMB's focus on promoting strong risk management practices and removing barriers to the use of AI by federal agencies. We highlight three important aspects of this part of the OMB memo.

- First, the OMB memo's requirement for agencies to designate a Chief AI Officer is an important aspect of implementing AI risk management programs across agencies. Designating specific personnel responsible for addressing AI is a key part of developing a successful AI governance program. The Chief AI Officer can fulfill important functions to help promote agencies' use of AI and manage AI risks.
- Second, we appreciate the OMB memo's focus on removing barriers to the use of AI. The OMB memo identifies key mechanisms for achieving this goal, including ensuring access to adequate IT infrastructure and encouraging the use of agency and public access datasets. We encourage agencies to ensure access to multi-cloud environments as part of their IT modernization. We also encourage federal agencies to expand open and secure access to federal data, recognizing that government-generated data is an important asset that can serve as a powerful engine for creating new jobs, promoting economic growth, and enabling innovation in AI-related technologies. Agencies collect and generate vast quantities of data that offer unique insights into virtually every facet of the modern world, from

satellite imagery that can help predict the weather to transportation data that can help reduce congestion. To enhance AI innovation, agencies should continue to prioritize the release of high-value, non-sensitive government data.

- Third, the OMB memo aims to reduce barriers to the use of AI by filling gaps in AI talent. The focus on increasing and training the workforce for AI positions will be a critical part of advancing AI innovation across the federal government. Global competition to train and recruit the next generation of tech talent is fierce, and solving the high-tech workforce shortage will require a broad strategy including the practices identified in the OMB memo for recruiting, training, and retaining workers.

## **II. The OMB Memo Should Ensure Uniform Application of the Definitions of Rights-Impacting and Safety-Impacting AI Across Federal Agencies and Focus on Risks That Impact Individuals.**

The OMB memo requires agencies to implement minimum risk management practices for rights-impacting and safety-impacting AI. It also includes a list of purposes for which AI is presumed to be rights-impacting or safety-impacting, while encouraging agencies to supplement this list by adding additional contexts where AI tools would be subject to heightened requirements. We recommend revising the OMB memo to address four issues raised by the memo's approach to defining rights-impacting and safety-impacting AI.

- First, at the outset, we note that the definitions in Section 6 of rights-impacting and safety-impacting AI are overbroad and ambiguous. For example, the definition of safety-impacting AI refers to “AI that has the potential to meaningfully impact” safety, and the definition of rights-impacting AI refers to “non-deceptive information about goods and services.” The reference to “potential” unnecessarily expands the scope of the definition to scenarios that don't actually pose safety risks, just as the reference to information about goods and services could inadvertently and significantly expand the scope of this language. In lieu of applying ambiguous definitions of safety-impacting and rights-impacting AI, **OMB should focus on high-risk use cases resulting in consequential decisions about individuals, namely decisions that create legal or similarly significant effects on individuals.**
- Second, to the extent that OMB maintains these categories, including a non-exhaustive list of rights-impacting and safety-impacting AI in the OMB memo while allowing agencies to develop additional agency-specific lists creates fragmentation. This undermines the US government's ability to develop a uniform government-wide approach to AI risk management that applies consistent standards across agencies. Under the current approach, two different agencies could evaluate the same AI system and reach conflicting results about whether it is rights-impacting or safety-impacting. This leads to inconsistent protection from AI risks. It also creates unpredictability for vendors, who will not know whether their AI system will be treated in the same manner across federal agencies — or if agencies will apply different obligations for the same AI system. A fragmented approach to AI governance also raises costs for agencies because it leads to duplicative assessments and because vendors may not be able to provide the same software services across different agencies. **We recommend revising the OMB memo to include an exhaustive list of purposes and use cases that are rights-impacting and safety-impacting.** Alternatively, OMB could be provided the authority to update this list over time, at set intervals, which would create a more stable approach than creating an agency-by-agency approach.

- Third, the threshold for determining whether AI is rights-impacting or safety-impacting is vague and can lead to differing interpretations across agencies. This heightens the risks of fragmentation discussed above. Specifically, the OMB memo states that AI is rights-impacting or safety-impacting “if it is used to control or meaningfully influence the outcomes” of the enumerated activities. The phrase “meaningfully influence” is susceptible to varied interpretations and does not provide a clear threshold for identifying instances in which an outcome is the direct result of AI. **We recommend omitting the phrase “meaningfully influence” and relying solely on the concept of “control.”**
- Fourth, the list of AI that is presumed to be rights-impacting and safety-impacting captures a broad set of AI applications, rather than specific use cases. This expansive list of applications and sectors would ultimately cover beneficial use cases, putting innovation at risk in very dynamic areas, such as cybersecurity. The approach to employment is also overly broad, lacking a clear focus on specific uses that create heightened risks. For example, the OMB memo references virtual workplace training programs, although it is not clear how or why these uses of AI would create meaningful risks to employees. Key terms such as “performance management” are also undefined. **We recommend that OMB tailor this section to focus on consequential employment decisions – decisions by agencies to hire, terminate, determine the pay, or promote individuals.**

### III. The OMB Memo Should Better Leverage the NIST AI RMF.

We urge OMB to revise the memo to better leverage the NIST AI RMF.

The OMB memo currently refers to the NIST AI RMF as a gap-filling tool. It encourages agencies to identify additional context-specific risks and to fill potential risk management gaps by incorporating the NIST AI RMF, the Blueprint for an AI Bill of Rights, applicable international standards, or the workforce principles established pursuant to the AI executive order. This approach treats the NIST AI RMF as one of many reference points, without looking to the NIST AI RMF as a key aspect of implementing the memo’s required minimum risk management practices. We note that, in contrast to the policies, principles, or international standards referenced, the NIST AI RMF was developed at the direction of Congress.

We encourage OMB to better leverage the NIST AI RMF, which is a flexible framework for identifying and mitigating risks. Adopting the AI RMF across government agencies would make the US a market leader on responsible AI and embrace best practices for managing AI risks. Leveraging the AI RMF can also help address the memo’s fragmented approach to AI governance across the federal government, because the AI RMF can serve as a common language for assessing agencies’ AI governance activities.

Creating the NIST AI RMF was a significant achievement that builds on NIST’s work creating frameworks for managing cybersecurity and privacy risks. For example, the NIST Cybersecurity Framework (CSF) is widely used by private and public-sector organizations worldwide; since 2017, it has been mandatory for federal agencies to use the CSF to improve their cybersecurity risk management programs.<sup>5</sup> **We strongly encourage OMB to**

---

<sup>5</sup> 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>.

**incorporate the NIST AI RMF as part of the memo's minimum risk management practices for federal agencies.**

#### **IV. The OMB Memo Should Refine its Minimum Risk Management Practices.**

The OMB memo requires federal agencies to adopt minimum risk management practices for rights-impacting and safety-impacting AI. These requirements are designed to help agencies manage AI risks. We offer several recommendations on how OMB can best support agencies' implementation of these practices.

##### **A. Impact assessments**

The OMB memo recognizes the importance of performing impact assessments for high-risk AI systems, by requiring federal agencies to complete an impact assessment for both rights-impacting AI and safety-impacting AI. Performing impact assessments is a key part of creating a strong 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.<sup>6</sup> We appreciate that the OMB memo recognizes the utility of this important accountability tool and offer three recommendations as you implement this obligation:

- **First, the OMB memo should be revised to recognize that an agency's impact assessment should reflect the agency's role in obtaining and using the AI system at issue.** This is important because different agencies will play different roles in developing and using AI systems. As a result, they will have access to different types of information and be positioned to take different types of actions to identify and mitigate risks associated with the use of an AI system. We strongly recommend revising the OMB guidance to incorporate the overarching point that the impact assessment should reflect the agency's role in developing or using the AI system at issue.
- **Second, the OMB memo should not require agencies to provide specific metrics or qualitative analysis supporting its intended purpose and expected benefits.** The OMB memo currently requires an impact assessment to specify not just the intended purpose of the AI system, but also quantifiable metrics or qualitative analysis of the intended purpose and its expected benefits. This obligation should be omitted because it is likely to create barriers to using AI, which is in tension with the stated goal of removing barriers to the responsible use of AI.
- **Third, the OMB memo should streamline its requirements for documenting data used in the development and testing of AI.** The OMB memo requires documentation of several items, including the provenance and quality of data for its intended purpose; whether the data comes from an adequately reliable source; and whether the data is sufficiently broad enough to address the real-world inputs it might encounter. However, to the extent agencies rely on vendors for AI systems, the disclosure of this information is likely to implicate proprietary information and create a barrier to using vendor AI unless the relevant information has already been made available to the public or the downstream provider with whom the agency has contracted. It could also unintentionally undermine privacy and security of the AI

---

<sup>6</sup> See BSA, Impact Assessments: A Key Part of AI Accountability, *available at* <https://www.bsa.org/files/policy-filings/08012023impactassess.pdf>.

system. For example, disclosures that require information detailing how network defenders use and train AI systems to secure networks could unintentionally create a roadmap for cyber adversaries to break through those defenses, in turn jeopardizing the underlying security of network and information systems. To address this concern, OMB should adjust these requirements to make them more high-level, such as requiring an overview of the training data.

- **Fourth, the OMB memo should provide that an agency can rely on another agency's impact assessment if the use case and context are the same or substantially the same.** This approach both enhances efficiency and promotes consistency among agencies' implementation.
- **Fifth, OMB should set a time frame for agencies to carry out impact assessment requirements for those conducted after the initial implementation deadline.** This will ensure that the implementation of safeguards does not slow down the adoption of AI.

## **B. Identifying algorithmic discrimination**

The OMB memo requires agencies to take steps to ensure AI will advance equity, dignity, and fairness, as an additional obligation for rights-impacting AI.

The goal of this provision is important, as the need to identify and mitigate potential biases in AI systems is a critical issue. 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 unlawful bias. We also believe that civil rights laws should be fully enforced.

At the same time, we encourage OMB to revise this aspect of the memo, to recognize several practical challenges that arise in identifying and mitigating potential biases. For example, the OMB memo currently requires agencies to identify and remove factors contributing to algorithmic discrimination or bias for rights-impacting AI. This involves determining if the AI relies on information about a protected class under federal nondiscrimination laws. However, there may be a lack of information for agencies to make these assessments, such as when the training data at issue does not contain information about protected classes. In addition, agencies will have to balance their privacy obligations, which can limit their ability to collect additional information that may provide insight on potential biases, with the need to assess AI for bias and discrimination.<sup>7</sup> **We urge OMB to further consult with stakeholders on the best approaches for agencies to identify potential biases and create flexibility in the memo by recognizing that agencies may not have access to information that would assist them in identifying potential biases.**

## **C. Consultation with stakeholders**

The OMB memo also requires agencies to consult and incorporate feedback from affected groups, as an additional minimum practice for rights-impacting AI. Although we recognize

---

<sup>7</sup> See Jennifer King, Daniel Ho, et al., The Privacy-Bias Tradeoff: Data Minimization and Racial Disparity Assessments in U.S. Government, *available at* <https://dl.acm.org/doi/10.1145/3593013.3594015>.

the importance of stakeholder feedback, we urge OMB to revise the memo's approach to soliciting and responding to such feedback in this context.

The OMB memo requires OMB to consult affected groups in the development and use of rights-impacting AI. If agencies receive "negative feedback," OMB advises them to consider not deploying the AI or removing the AI from use. While it is important for stakeholders to provide input on an agency's policies and overarching approach to important issues, the memo focuses stakeholder feedback on how agencies use a specific AI product, as opposed to a broader policy approach to AI technologies. This product-specific approach creates the potential for stakeholders to veto or otherwise slow down the use of a specific AI system. The process is also open to abuse by competitors and may become a second forum for bid protests. Moreover, the standard for "negative feedback" is vague and undefined. **OMB should revise the memo to advise agencies to take feedback into account without encouraging them to terminate use of a particular AI product in the event of negative feedback.**

#### V. The OMB Memo Should Extend the Timeline for Agency Implementation.

The OMB memo requires agencies to implement the minimum risk management practices and terminate non-compliant AI by August 1, 2024. Given the complexity of the issues that agencies are addressing, the implementation deadline is too soon for agencies to implement the required practices in advance of the stated deadline.

We strongly recommend agencies be given more time. If this deadline is not extended, agencies risk ending up in a scenario where they are not able to use AI simply because they have not been able to implement a substantial set of new practices across the agency within this short time frame. **We recognize that agencies can seek extensions under the OMB memo, but agency-by-agency extensions should not be a substitute for establishing an initial deadline that sets a strong and workable timeline for all agencies to adopt and implement the required new practices.**

#### VI. OMB Should Revise the Memo's Procurement Recommendations.

The OMB memo also addresses the process by which federal agencies procure AI systems. We recommend revising the memo's procurement provisions in three ways:

- **First, OMB should revise the memo to establish standard practices across agencies.** The procurement section of the OMB memo currently takes an agency-by-agency approach to procurement. For example, it requires agencies to include "tailored risk management requirements" in procurement contracts, and to encourage individual agencies to obtain adequate documentation of procured AI. We recommend revising this approach, so that procurement requirements are standardized across agencies. A standard process can better permit federal agencies to procure technology quickly. In contrast, allowing agencies to independently set these procurement requirements means that companies must tailor their systems for each agency, which takes time and increases costs to the taxpayer. The OMB memo should create a standard set of government-wide practices for AI procurement, helping ensure agencies can effectively adopt the technology.
- **Second, OMB should revise the memo's approach to external testing of procured generative AI systems.** The OMB memo encourages agencies to incorporate tailored risk management requirements in contracts for generative AI.

These requirements include obligations for testing and safeguards, including external red teaming. We urge OMB to revise the memo's approach to these issues and delete the reference to "external" red teaming. Testing is an important part of ensuring appropriate functionality, but involving external entities in such testing can create concerns around access to trade secrets, proprietary information, or information that could jeopardize network and information security, and is likely to create barriers to using AI. Encouraging internal testing avoids requirements for companies to disclose this sensitive information to third parties and run the risk of such information ending up in the wrong hands.

- **Third, OMB should support the improvement of procured AI systems.** The OMB memo includes important provisions that encourage federal agencies to consider contracting provisions that incentivize the continuous improvement of procured AI. However, the memo also recommends that agencies consider contracting provisions that ensure that government data cannot be used to train or improve functionality of commercial AI offerings by the vendor without express permission from the agency. This creates tension with the broader goal of ensuring agencies have access to continually improved AI. Allowing vendors to use agencies' data to train their commercial products to continue to improve the technology is critical to ensuring the continuous improvement of procured AI, and to ensuring an AI system can more accurately support the agency's goals and data types. We recommend that OMB revise the memo to encourage agencies to permit their data to be used for further training of AI systems used by the government, subject to appropriate privacy and security safeguards.

## **VII. The OMB Memo Must Be Harmonized With Concurrent Regulatory Actions.**

The OMB memo includes recommendations to agencies for AI procurement. We strongly encourage OMB to ensure the memo takes account of other ongoing regulatory actions that intersect with issues at the center of the OMB memo.

We recommend the memo acknowledge the need for OMB to work towards harmonizing the range of draft regulations and guidance that are currently out for industry comment and review. These include the OMB's memo on FedRAMP, "Modernizing the Federal Risk Authorization Management Program." The extension for comments on that draft guidance is helpful as the FedRAMP regulations will impact the larger cloud market, which in turn, will impact the AI market.<sup>8</sup> Moreover, the Federal Acquisitions Council is considering three rules for which the comment period was extended to February 2, 2024 (FAR-2021-0017 Federal Acquisition Regulation: Cyber Threat and Incident Reporting and Information Sharing; FAR-2021-0019 Federal Acquisition Regulation: Standardizing Cybersecurity Requirements for Unclassified Federal Information Systems), and (FAR- 2023 – 06 Federal Acquisition Regulation: Implementation of Federal Acquisition Supply Chain Security Act (FASCASA) Orders).

**We urge OMB to consult with stakeholders including industry to look at these changes in concert, as the multiple regulations will need a coordinated approach to advancing AI innovation.** With the number of changes in the regulatory environments for IT and

---

<sup>8</sup> These issues are closely related, as different AI models may have stronger performance in different clouds, making competition in the cloud marketplace crucial.



specifically, AI, thoughtful analysis on the interplay between these rules and regulations can help the federal government better achieve its stated goals.

\* \* \*

We appreciate the opportunity to provide comments on the OMB memo and would be happy to serve as a resource as you continue to develop your approach to these issues.