Dec. 5, 2023

The Honorable Shalanda Young
Director, Office of Management and Budget
725 7th Street NW
Washington, DC 20503

Re: Notice of Public Comment Period, Office of Management and Budget; Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (88 FR 75625, OMB-2023-0020), Nov. 3, 2023

Dear Director Young:

The Security Industry Association (SIA) appreciates the opportunity to provide comments to the Office of Management and Budget (OMB) on its draft memorandum, "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence" ("draft memorandum").[1] The draft memorandum, among other things, would require agencies to apply a "minimum baseline" of practices to manage risks of "rights-impacting AI" and "safety-impacting AI." OMB also includes a series of recommendations to apply risk management approaches to federal procurement and intends to develop a system for ensuring federal contracts align with the Oct. 30, 2023, Executive Order for Safe Secure and Trustworthy Artificial Intelligence ("AI EO").

About SIA

SIA represents more than 1,400 safety and security solutions providers, ranging from large global technology firms to locally owned and operated small businesses. Our members include manufacturers, software developers and systems integrators providing solutions that help protect people, property and sensitive data across a wide variety of public- and private-sector applications.

We believe advanced technologies must only be used in a lawful, ethical, nondiscriminatory and effective manner. SIA and our members support the development and use of artificial intelligence (AI) technologies in ways that are human-centric, ethical and trustworthy and that mitigate potential risks. In part to help communicate and guide effective technology implementations that harness these innovations, SIA has formed its AI Advisory Board[2] and Identity and Biometric Technology Advisory

[1] See - https://ai.gov/wp-content/uploads/2023/11/AI-in-Government-Memo-Public-Comment.pdf
[2] See - https://www.securityindustry.org/committee/ai-advisory-board/

Board.[3] SIA is also pleased to have a representative serving on the National AI Advisory Committee Subcommittee on Law Enforcement.[4]

AI-Driven Safety and Security Technologies Are Critical to Agency Missions

AI stands to provide enormous individual, societal and economic benefits in almost every sector. Within the safety and security space, AI helps to power advanced tools, enabling technology users to analyze and respond to information in a substantially quicker and more accurate manner than traditional, non-automated methods. For example, AI tools and technologies can streamline 911 call center responses, better transcribe incident reports, and analyze video feeds for high-risk safety and security situations at faster speeds. Associated techniques like machine learning, computer vision, deep learning and neural networks can dramatically improve the capabilities of security solutions and have been the primary areas of recent technological innovation in our industry, which contributes over $150 billion to the economy and employs more than a million people across the United States.

Deploying these technologies results in more preventative, as opposed to reactionary, capabilities for the security of government facilities and greater effectiveness in the performance of critical agency missions related to homeland security, law enforcement, public safety, emergency response, identity management, cybersecurity, program integrity measures and many more. Importantly, AI can also help protect against some of the broader potential risks that OMB's proposed policies are intended to prevent. For example, there is a growing body of research that shows how AI can identify and mitigate against bias in human decision making, thereby moving towards the goal of preventing unlawful bias.[5] And effective AI deployment can also help improve safety and security.

Comment on the Draft Memorandum

The draft memorandum follows and specifically responds to Section 10.1(b) of the AI EO, requiring OMB to, within 150 days, "issue guidance to agencies to strengthen the effective and appropriate use of AI, advance AI innovation, and manage risks from AI in the Federal Government."

We fully support the aims of the draft memorandum, as well as the eight principles the AI EO directs federal agencies to adhere to in the development and use of AI: 1) ensuring AI safety and security, 2)

---

[3] See - https://www.securityindustry.org/committee/identity-and-biometric-technology-advisory-board-ibtab/

[4] See - https://www.ai.gov/naiac/#SUBCOMMITTEE_ON_AI_AND_LAW_ENFORCEMENT_NAIAC-LE

[5] *See e.g., Mitigation of AI/ML Bias in Context*, NCCoE, https://www.nccoe.nist.gov/projects/mitigating-aiml-bias-context ("Automated decision-making is appealing because artificial intelligence (AI)/machine learning (ML) systems produce more consistent, traceable, and repeatable decisions compared to humans . . . ."); Kimberly Houser, *Can AI Solve the Diversity Problem in the Tech Industry: Mitigating Noise and Bias in Employment Decision-Making*, 22 Stan. Tech. L. Rev. 290, 352 (2019), https://www-cdn.law.stanford.edu/wp-content/uploads/2019/08/Houser_20190830_test.pdf ("The use of AI in talent-management decisions has shown success in not only creating more successful hires, but in also creating a more diverse slate of candidates and employees. While some companies have embraced these new technologies, others fear that AI may actually cause discriminatory outcomes. As discussed, the phenomena of 'garbage in, garbage out' is real, but can be addressed paying attention to the data sets by using known sources and making sure the sets are balanced and representative of all groups."); Jon Kleinberg et al., *Discrimination in the Age of Algorithms*, 10 J. of Legal Analysis 113, 120 (Apr. 22, 2019), https://academic.oup.com/jla/article/doi/10.1093/jla/laz001/5476086 ("Our central claim, stated in simple form, is that safeguards against the biases of the people who build algorithms, rather than against algorithms per se, could play a key role in ensuring that algorithms are not being built in a way that discriminates (recognizing the complexity and contested character of that term). If we do that, then algorithms go beyond merely being a threat to be regulated; they can also be a positive force for social justice.").

promoting responsible innovation, competition and collaboration, 3) supporting American workers, 4) advancing equity and civil rights, 5) standing up for consumers, patients and students, 6) protecting Americans' privacy and civil liberties, 7) ensuring responsible and effective government use of AI and 8) advancing American leadership abroad.

That said, we are concerned that, without changes, there are ***significant issues in the draft memorandum that could lead to results that undermine these goals and a true risk-based approach***. Most notably, we are concerned that the draft memorandum's broad definitions of "rights-impacting AI" and "safety-impacting AI" combined with a broad and rigid checklist of "minimum practices" associated with their use will result in discouraging the use of AI in the federal government. Further, given the limited resources of federal agencies, these broad definitions and one-size-fits-all requirements could foster an environment in which truly high-risk uses of AI do not receive proportional attention compared to lower-risk AI uses that happen to fall under the definitions of "rights-impacting" and "safety-impacting" AI. We offer the following recommendations in three key areas to address these issues in finalizing the memorandum.

1. **AI use cases subject to "minimum requirements" must be narrowed and clarified to ensure the framework appropriately focuses on applications that are truly high-risk.**

The draft memorandum is framed as suggesting a targeted approach to AI use by agencies, in that it only proposes to impose minimum requirements on "rights-impacting AI" and "safety-impacting AI." However, it broadly defines "rights-impacting AI"[6] and "safety-impacting AI,"[7] and additionally lists broad categories of use cases that are automatically presumed to fall under these definitions. This overly broad approach risks categorizing nearly *all* AI use cases with any connection to the listed categories as high-risk, which would undermine the goals of the AI EO and the OMB memorandum, and stifle AI use and development to support federal agency missions.

In particular, the draft memorandum's sweeping definitions do not adequately differentiate between different levels of risk. The definition of "rights-impacting AI" includes, for example, "AI whose output serves as a basis for a decision or action that has a legal, material, or similarly significant effect" on individuals' or communities' privacy and access to critical resources or services, among other things.[8] Without adequate tailoring or guidance, these concepts can be construed broadly and sweep in a vast

[6] "AI whose output serves as a basis for decision or action that has a legal, material, or similarly significant effect on an individual's or community's: 1. Civil rights, civil liberties, or privacy, including but not limited to freedom of speech, voting, human autonomy, and protections from discrimination, excessive punishment, and unlawful surveillance; 2. Equal opportunities, including equitable access to education, housing, credit, employment, and other programs where civil rights and equal opportunity protections apply; or 3. Access to critical resources or services, including healthcare, financial services, social services, transportation, non-deceptive information about goods and services, and government benefits or privileges." Draft memorandum at 24.

[7] "AI that has the potential to meaningfully impact the safety of: 1. Human life or well-being, including loss of life, serious injury, bodily harm, biological or chemical harms, occupational hazards, harassment or abuse, or mental health, including both individual and community aspects of these harms; 2. Climate or environment, including irreversible or significant environmental damage; 3. Critical infrastructure, including the critical infrastructure sectors defined in Presidential Policy Directive 2143 and the infrastructure for voting and protecting the integrity of elections; or, 4. Strategic assets or resources, including high-value property, information marked as sensitive or classified by the Federal Government, and intellectual property." Draft memorandum at 25.

[8] Draft memorandum at 24.

array of AI use cases – for example, AI outputs that are just one factor in a decision made by an individual, and/or that have a potential effect on "privacy" generally even if there is no directly applicable privacy law.

The definition of "safety-impacting AI" is similarly broad. This category includes AI that "has the potential to meaningfully impact the safety" of a range of things, from "human life or well-being" to "critical infrastructure" and "[s]trategic assets or resources." The inclusion of the word "potential", as well as the broad and vague range of impacts – for example, "well-being" – does not provide meaningful parameters to help federal agencies draw lines between high-risk applications and other applications. This approach is likely to result in a set of policies and practices that may ultimately fall short of accomplishing the intended risk-based approach to regulation because it does not sufficiently differentiate, at a practical level, between that which may actually cause harm and that which is much less likely to cause harm.

Additionally, the lists of specific technologies and use cases that the draft memorandum "presume[s] to be rights-impacting" or "presume[s] to be safety-impacting" are overly broad and do not fully account for the wide variety of use cases of AI applications within the established categories. For instance, the list for "rights-impacting" AI includes broad categories such as "facial matching," "physical location-monitoring devices," and "license plate readers."[9] While such technologies – already utilized for two decades in some cases – could be high-risk under certain contexts, grouping all instances of these AI use cases together without differentiation between different risk levels is counter to a risk-based approach and could result in agencies being discouraged from using these applications for beneficial purposes even when the agency might be able to do so in a sufficiently narrow and low-risk manner.

For example, license-plate readers and plate data can be used to provide a number of services beneficial to the public that would be difficult to construe as "high-risk" such as traffic analysis and management, toll automation and smart cities applications, such as intelligent parking. Moreover, the specific AI utilized in such applications is limited to a narrow portion of the application workflow – utilizing computer vision for the purpose of character recognition. Due to the nature of this use case and the narrow capabilities of the supporting AI, it is difficult to conceive a path for the AI in this use case to provide any substantial social or security risk. The draft approach to regulating AI set forth in the memorandum, however, does not provide a sufficient framework to differentiate such a lower-risk use case from a higher-risk use case. Blanketly categorizing technologies as high-risk would undermine OMB's goals to advance AI innovation and "seize the opportunities AI presents,"[10] including to improve safety and security across the federal government.

The technologies laid out as presumptively "safety-impacting" in the draft memorandum also cover broad categories of applications that – depending on the context of deployment – may or may not be high-risk. For example, the list includes "[r]esponses to insider threats," "[a]ccess to and security of government facilities," and "movements of vehicles."[11] Again, including these blanket examples as presumptively "safety-impacting" could negatively impact the federal government's effective use of AI for beneficial purposes, including to assist it with cybersecurity and defense activities.

[9] *Id.* at 12.
[10] *Id.* at 1.
[11] *Id.* at 11.

Additionally, the draft memorandum establishes an overly broad and vague standard for when these categories trigger the presumption, explaining that a specific use case is presumptively rights- or safety-impacting when AI "is used to control or meaningfully influence the outcomes of the [listed categories of activities]." The standard of "meaningfully influenc[ing] the outcome" is vague and difficult to operationalize and risks being construed broadly to potentially include any operation that includes AI – which would be in direct tension with a risk management approach. Instead, the final OMB memorandum should, 1) strike "the outcomes of" language for clarity (so the standard would be if "it is used to control or meaningfully influence the following activities") and 2) provide clear guidance on the parameters or outer limits of what it means to "meaningfully influence." At a minimum, the memorandum should make clear that the following activities would not meet the "meaningfully influence" standard, consistent with recent proposals for high-risk categorization under consideration in the European Union's AI Act, which exclude when: 1) "the AI system is intended to perform a narrow procedural task of low complexity," 2) "the use of the AI system does not replace a human assessment relevant for the purpose of the use case [at issue], but is intended to confirm or improve an accessory factor of such assessment;" or 3) "the AI system is intended to perform a preparatory task to an assessment relevant for the purpose of the use case [at issue]."[12]

In short, we recommend that definitions for "rights-impacting" and "safety-impacting" in the draft memorandum should be narrowed to focus on truly high-risk applications, rather than listing broad categories of technologies *presumptively* falling into its high-risk framework. In addition, guidance should be provided to agencies to help assess the unique set of benefits and risks to a specific AI application, through objective and not subjective standards,[13] before classifying it as either "rights-impacting" or "safety-impacting."

2. **Agencies should be provided with more flexibility regarding "minimum practices" to tailor risk mitigation strategies to specific AI applications in greater alignment with the National Institute of Standards and Technology (NIST) Risk Management Framework (RMF).**

AI encompasses a wide range of emerging technologies with countless iterations and applications. There is no "one-size-fits-all" approach that can be used to mitigate risks across all AI applications as the technology evolves. As currently written, the draft memorandum creates a rigid checklist applicable to "rights-impacting AI" or "safety-impacting AI" that is too prescriptive. Flexibility is a longstanding and important risk management principle that addresses this issue: approaches to AI should be flexible enough to account for the wide range of ways that AI can be deployed for beneficial uses and a range of methods to appropriately address risks.

[12] See https://www.notizie.ai/pathal/uploads/2023/10/20231003-ai-act-article-6_new-filter.pdf

[13] As an analogy, the U.K. Information Commissioner's Office provides objective guidance on when an activity is presumed "high risk" (see - https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/accountability-and-governance/data-protection-impact-assessments-dpias/when-do-we-need-to-do-a-dpia/). For instance: not generically presuming the processing of sensitive data is high risk, but rather the processing of sensitive data on a large scale. Creating objective standards need not undermine each Agency's autonomy to assess (within their scope of practice and each specific use case) whether certain AI applications fall within the categories of safety-impacting or risk-impacting. By curtailing agencies autonomy to do so through generic pre-formulated categories, the draft memorandum undermines its own stated goal of allowing each Agency to conduct their own review, as the same use of a certain AI application by one agency might be different than that same use by another agency.

Unfortunately, as currently written, the draft memorandum does not afford such flexibility. Instead, it creates a rigid checklist applicable to "rights-impacting AI" or "safety-impacting AI" that is too prescriptive, particularly in light of the unsettled and rapidly evolving landscape of AI safety.

This approach stands in stark contrast to the risk-based and flexible AI RMF, which establishes governance principles and goals for AI risk management but does not include specific methods for implementation.[14] Published in January 2023, the AI RMF, together with its accompanying AI RMF Playbook,[15] is a tool for organizations of varying types and sizes to identify and manage benefits and risks when developing and deploying AI technologies. During its development, NIST collaborated with a broad range of stakeholders, so the AI RMF is aligned with existing best practices and industry standards.[16] Importantly, it is meant to be applied flexibly to a variety of use cases. Since the second draft of the AI RMF, NIST has specifically noted that the "AI RMF is not a checklist."[17]

As it finalizes its memorandum, OMB should align its "minimum practices" requirements with NIST's AI RMF to ensure that federal agencies balance AI benefits and risks. Already, the AI EO[18] and the draft memorandum[19] point to and rely on the AI RMF as a standard for proper AI risk management. In fact, the AI EO explicitly instructs OMB to specify in its guidance "required minimum risk-management practices for government uses of AI that impact people's rights or safety, including, where appropriate, the following practices derived from…the NIST AI Risk Management Framework."[20] Building from the AI EO, OMB should make targeted updates to the draft memorandum to ensure that NIST's AI RMF is touchstone for OMB in crafting its AI risk mitigation requirements for federal agencies.

Specifically, to better align with the NIST RMF, OMB's guidance should provide a menu of options[21] for federal agencies to select from to mitigate risks. Agencies should perform a risk assessment and select appropriate controls – tailored to the context of the AI deployment – to best address specific risks; they should not have to apply a pre-determined checklist of detailed steps, no matter the context. By applying a more flexible, risk-based approach, consistent with the AI RMF, OMB can more effectively mitigate risks by allowing federal agencies to tailor their risk mitigation strategies to the actual risks their

[14] *See e.g.,* NIST AI RMF Playbook § Govern 1.2, https://airc.nist.gov/docs/AI_RMF_Playbook.pdf (NIST includes governance standards such as inclusion of "characteristics of trustworthy AI" in "organizational policies, processes, and procedures," but does not dictate how that is accomplished. Instead, the NIST AI RMF Playbook suggests actions such as defining and scoping key AI terms and concepts, conducting model testing and validation processes, and outlining change management requirements."

[15] https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf; https://airc.nist.gov/AI_RMF_Knowledge_Base/Playbook

[16] https://airc.nist.gov/AI_RMF_Knowledge_Base/Roadmap

[17] https://www.nist.gov/system/files/documents/2022/08/18/AI_RMF_2nd_draft.pdf.

[18] *See e.g.*, AI Executive Order at 75199.

[19] Draft memorandum at 13 ("To fill potential risk management gaps, agencies are encouraged to promote and to incorporate, as appropriate, additional best practices for AI risk management, such as from the National Institute of Standards and Technology (NIST) AI Risk Management Framework . . .").

[20] AI Executive Order at 75218.

[21] *See e.g.,* NIST AI RMF Playbook § Govern 1.3, https://airc.nist.gov/docs/AI_RMF_Playbook.pdf (The governance principle requires, "[p]rocesses and procedures are in place to determine the needed level of risk management activities based on the organization's risk tolerance," but the section subsequently lists five examples, "suggested actions," of types of policies that could be implemented to execute those types of risk management activities."); s*ee also supra* note 31 (providing options for implementation of a governing principle).

unique deployments face. This approach will also allow agencies to continue to adapt their approaches to prevailing best practices in AI safety as the field continues to evolve in the months and years ahead.

3. **Key adjustments are needed to "minimum practices" related to testing and disclosure of data quality and use in design, development and training.**

The draft memorandum's AI impact assessment requirements include data documentation and training data disclosures that either apply directly or flow down to technology providers.[22] We are supportive of ensuring the quality and appropriateness of data used in model development to help determine a model's performance for a particular use case; however, visibility into training and evaluation datasets will not meaningfully help an agency make that determination. The fact that a model has been trained on certain data does not directly indicate how it will perform in a specific use case. Performance testing, already required under the draft memorandum's "Minimum Practices," is the scientifically appropriate and relevant way to assess this.

Rather than focusing on training data, ensuring greater visibility into how a model works, in addition to its performance, would be a more effective approach. This could include requiring documentation regarding the capabilities of the AI system, known limitations, guidelines for use and demonstrative performance results.

The collection of training data used to develop an AI system or software is often intellectual property (IP) with significant competitive value, and disclosure could affect a provider's claim to trade secret protection, compromise proprietary information and data security (where disclosure could expose it to bad actors and possibly lead to data breaches), or lead to bias and misinterpretation (where disclosing the source of data might lead to erroneous conclusions about the performance of the AI application based on only its origin from a particular group, population or data set, potentially giving one technology provider an advantage over another based not on the performance of the AI application, but on the pre-judgement of its training data). This makes training data disclosure impractical and ill suited for agency evaluation. Therefore, any information requirement regarding training data should be limited to summary information about the data the model was trained on or that demonstrates how the training data is appropriate for the agency's intended use.

Additionally, we encourage OMB to provide direction to agencies, as they implement the Minimum Practices, that they should, to the extent possible, limit administrative duplication and potential unnecessary disruptions or delays in needed AI technology implementations, by incorporating existing relevant agency polices and documentation into the AI impact assessment. For example, many "safety-impacting" and "rights-impacting" AI systems will already be accompanied by extensive privacy impact assessments (PIAs) that include required elements. Also, some AI-driven technologies are already independently and robustly tested through government evaluation programs.[23]

[22] Draft memorandum at 15-16

[23] For example, face recognition and related technologies are evaluated by National Institute of Standards and Technology (NIST), https://www.nist.gov/programs-projects/face-recognition-vendor-test-frvt, and Department of Homeland Security, https://www.dhs.gov/science-and-technology/BI-TC.

Finally, regarding the practice outlined to "Consult and incorporate feedback from affected groups," we strongly encourage OMB to provide additional direction to agencies to engage with all relevant stakeholders – including technology providers and other subject matter experts – regarding "negative feedback" regarding specific AI systems, to help ensure such feedback has merit.

Conclusion

Agency requirements under the memorandum need to balance promoting beneficial uses of AI while managing risks. The Request for Comment asks how "OMB [can] best advance responsible AI innovation."[24] The draft memorandum recognizes that AI can "improve operations and deliver efficiencies across the Federal Government"[25] and encourages agencies to identify the opportunities and benefits AI can provide to the agencies mission and "improve their ability to use AI in ways that benefit the public and increase mission effectiveness."[26]

We are concerned that, without changes, the broad scope and overly prescriptive nature of certain requirements in the draft memorandum could put up roadblocks to the federal government using AI and potentially discourage companies with AI-powered tools and solutions from engaging with federal government.

Additionally, as polices are developed to ensure federal contract alignment with the AI EO, we urge OMB to facilitate informed decision making in the federal procurement process with clear and specific requirements for procuring AI products and services. OMB should also consider creating additional guidance or tools that would help current and potential suppliers assess compliance requirements in relations to their offerings ahead of time, which could create efficiencies in procurement processes and throughout contract life cycles.

Again, we appreciate the opportunity to provide initial feedback on the draft memorandum. SIA and its members stand ready to work with OMB to ensure that government can continue to leverage AI technologies in a safe, secure and effective manner.

Sincerely,

Don Erickson
Chief Executive Officer
Security Industry Association
Silver Spring, MD
www.securityindustry.org

Staff Contact: Jake Parker, jparker@securityindustry.org

[24] Draft memorandum RFC at Question 3.
[25] Draft memorandum at 8.
[26] *Id.*