Openpolicy

Jan. 5, 2024

To: The Office of Management and Budget (OMB), via ofcio@omb.eop.gov

Atten: Clare Martorana, Conrad Stosz

OpenPolicy comments on

OMB Draft Memo on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (AI) (the "Memo")

Overview

OpenPolicy appreciates the opportunity to provide feedback on the OMB Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (AI) Memo. OpenPolicy is a technology company seeking to democratize the ability of innovative companies of all sizes to engage with policymakers and provide feedback on relevant policy deliverables. OpenPolicy is further engaged in emphasizing the need to use AI, innovation, and technology to foster open policymaking and broader use of technology to streamline compliance and governance automation. We anticipate active engagement with OMB throughout the evolution of the proposed Memo and are available to provide further feedback as needed. We are committed to supporting the implementation of the Executive Order and actively engaging with OMB and other implementing agencies, alongside our participation in the NISA AI Safety Institute Consortium.

We believe that the open and collaborative nature of the policymaking dialogue is essential to further the implementation of the Executive Order, which can be significantly contributed by the participation of innovative companies and, specifically, startup companies that develop cutting-edge AI security and safety solutions. Indeed, many, if not most, of the technologies used to support the requirements of the AI EO and relevant OMB and the underlying NIST guidelines referred to in the Memo (such as NIST RMF) evaluate the measurements, testing, and audibility of AI, data and security posture, and facilitate the secure adoption of AI and sharing of data, more broadly, are developed by such innovative companies and startups – these are the communities OpenPolicy collaborates with.

Moving forward, we recommend extending the comment period for key deliverables such as this one to ensure an opportunity for ample public participation and input and extend additional efforts such as roundtables and active outreach to ensure the voice of such communities is heard.

OpenPolicy applauds OMB for its leadership in creating and developing the Memo guidelines, outlining direct responsibilities and best practices, and, more broadly, efforts to promote and introduce new agency mandates concerning AI and data governance, innovation, and risk management. Guiding agencies in implementing specific minimum risk management practices

for AI applications affecting public rights and safety is critical for protecting national and federal systems. Additionally, the draft guidance focuses on removing unnecessary barriers to responsible AI use, including IT infrastructure limitations, dataset accessibility issues, cybersecurity authorization hurdles, and AI talent shortages, which are all key areas where substantial advancement is needed to advance the administration's goals.

We are especially satisfied with the emphasis on governance and automation, fostering enhanced measurability of AI security and safety outcomes and acknowledging the interconnected nature of AI with data, and the need to facilitate secure and safe data sharing; many of these measurement approaches can be connected to other government initiatives, such as the FedRamp revisions, to use of tools like OSCAL for ensuring control mappings as more AI security controls are added, and the on-going implementation of the Cyber EO.

Nevertheless, OpenPolicy encourages OMB to explore additional opportunities to leverage technology and automation tools to ensure minimum practices align with advanced responsible AI behaviors and procedures as they develop.

The current approach of processes related to measurements, alongside the establishment of new official roles, we fear, may not be sufficient, absent additional measures, to enhance Agency posture measurements and compliance efforts, as years of experience from security showcase.

More robust measurements scaled by automation and tech monitoring, alongside third-party tooling, are needed to scale compliance, and compliance measurements should be coupled with robust accountability measures to ensure agency preparation, alongside, of course, ample budget and workforce.

Such automation and measurement tolling are key to scale risk management, governance, and assessment methods to keep up to date with controls released and enable OMB to monitor agency implementation best. Automation can further support the complex goal of aligning the implementation of the Cyber Executive Order with the ongoing implementation of the AI Executive Order and upcoming revisions of FedRamp FISMA, FISMA 2024 priorities related to IoT and OT protections (where AI can be leveraged), and other NIST documents. A further focus on leveraging automation and measurable compliance (e.g., schemes like OSCAL) can advance risk governance in this respect and increase the ability of OMB to scale agency measurements, as agencies and auditors can collaborate with solutions providers to automate control measurements. Our ecosystem of innovative companies specializing in AI security measurements is eager to work with the administration to explore how technology and automation can best support the deployment of AI, security, and privacy-related controls and their effective enforcement and measurement.

In addition, we are delighted to offer the following specific comments on the Memo:

- **Alignment with other ongoing initiatives**

The implementation of the Memo requirements should align with upcoming revisions to NIST guidelines and controls and recent agency (e.g., CISA) published best practices and foster further coherence between the Cyber EO (14028), AI EO, and related efforts such as CMMC) controls' deployment by OMB and applicable National Security Systems enforcement efforts.

While OMB undergoes implementation of the AI EO, several key related documents that outline controls related to the Memo are being updated to comport with the threat landscape. Notably, the Cyber EO (14028) attestation form, currently under development by OMB, refers to NIST SSDF (SP 800-218), which is expected to be revised under the AI EO.

It is key that the Memo consider the need for coherence between the implementation of relevant cyber requirements and newly introduced AI cybersecurity requirements and controls under the ongoing implementation of the Cyber EO (14028) while we await final details of implementation related to AI EO, to ensure agencies and contractors are implementing the relevant controls, and not further cultivate adherence to past methods, that may expose the federal system to threats.

As an additional example, further consideration should be given to whether the Zero Trust Framework architecture by OMB needs revisions to align further with the newly developed AI requirements.

➢ Scope and strengthening AI Governance

1. *Scope*

    a. To effectively navigate the broad scope of requirements applied to AI and Data, the agency should consider a multi-faceted approach that leverages compliance measures, technology capabilities and testing, and automation within its operational framework to identify and mitigate risk. Addressing risks associated with AI use, as outlined in this memorandum, is pivotal. This includes leveraging risk governance and measurement technology measures, in particular, that can optimize and identify AI risks directly linked to the agency's AI applications and applicable data usage.

2. *Promoting AI Innovation*

    a. OMB's recommendations provide a comprehensive approach to fostering AI innovation within government agencies. By emphasizing the integration of AI into core mission objectives, removing barriers to adoption, and advocating for its benefits, agencies can effectively leverage AI to enhance their operations and services. The initiative to identify and prioritize appropriate AI applications aligns with the principles of responsible AI adoption. Agencies should carefully evaluate the potential impact of AI and the underlined data use cases and focus not only on applications that demonstrate clear

benefits and alignment with their strategic goals but also enable the use of AI to advance security and national security, to promote coherence and compliance with security controls and standards applicable more broadly. As emphasized in the Executive Order, the use of AI applications to advance security and identify vulnerabilities has a critical use case for AI in federal agencies.

b. The emphasis on removing barriers to AI adoption is crucial for accelerating its integration across government agencies. This includes investing in AI-enabling and data-sharing infrastructure, which is highly secured and includes the appropriate governance and audit controls. Agencies should also develop clear, supportive, and holistic policies that encourage responsible data sharing to enable AI technology adoption while providing access to necessary resources, such as funding and guidance, to ensure best practices, norms, behaviors, and data-sharing outcomes.

3. Managing Risks from the Use of AI

a. The recommendations outlined in the OMB Managing Risks from AI Use section effectively address the critical task of managing risks associated with AI adoption within government agencies. By establishing risk management programs, monitoring AI performance, ensuring compliance, conducting risk assessments by leveraging technology, and coordinating with relevant officials, agencies can effectively mitigate AI risks and ensure responsible AI practices. However, as the threat landscape evolves and additional security controls are developed by CISA, NIST, and implementing agencies, it is expected that the security controls for responsible agency adoption will evolve accordingly.

b. The initiative to establish or update risk management programs tailored explicitly to AI applications while conducting continuous monitoring and evaluation of AI applications is crucial for proactive risk mitigation and should further incorporate comprehensive risk assessment methodologies, clear risk categorization frameworks, and robust risk mitigation strategies that are supported by automation, technology capabilities, for compliance and posture management - to ensure effectiveness, safety, and adherence to controls. The provision for waiving individual AI applications from certain elements of the memorandum allows for flexibility in cases where strict adherence may not be feasible or may hinder innovation. However, rigorous justification and alternative risk mitigation measures should accompany this waiver process.

c. Rigorous compliance with AI-related requirements, including those established in the memorandum and other relevant laws and policies, is paramount for responsible AI adoption. Agencies should establish precise accountability mechanisms, differentiate the different components, and conduct regular audits to ensure adherence to these requirements. Further, developing agency-specific lists of purposes for which AI is

presumed to be safety-impacting, or rights-impacting should provide clarity and consistency in AI risk assessment and management and align with other agencies, such as the DOD and NIST, current requirements to ensure compliance and correlation with existing government mechanisms.

d. Collaboration among various implementing agencies and agency officials, including authorizing, procurement, legal, human capital, and oversight personnel, is essential for ensuring that AI adoption aligns with the memorandum's principles and risk management guidelines and vice versa. It is further essential that approaches to public-private partnerships and collaboration with the innovative ecosystem that operates on AI are considered as part of the memo development. This collaborative approach promotes holistic risk assessment and decision-making, ensuring compliance and secure AI adoption and management use. One proposal could allow OMB to leverage the NIST AI institute consortium to further public-private participation that can contribute to the Memo development and implementation.

➢ Advancing responsible AI innovation

1.AI Strategies

a. OMB's emphasis on secure data handling, storage, and disposal is essential. As agencies increasingly rely on AI, they must also take steps to protect the sensitive data that is used to train and operate AI systems, with governance and scaled audibility. This should include implementing robust cybersecurity measures, including automation and technology solutions to enhance the governance operation and function, reduce the risk of human error, and address emerging AI threat verticals stemming from the use of data, as well as AI red teaming. Automation for vulnerability management and contextualized risk assessment, alongside data posture management technologies, can be particularly beneficial for tasks such as data cleansing, data preparation, and model training.

b. The recommendation to establish mature AI-enabling infrastructure is also important and thus should include having the necessary computing resources, data storage capacity, and development tools in place. Developing and deploying AI systems effectively can be difficult without the proper infrastructure. This strategic approach would significantly benefit by encompassing secure data handling, storage, and disposal strategies, ultimately heightening protection for sensitive, government, and business-critical data and infrastructure. All of this requires the appropriate budgeting and workforce to deploy.

c. As the OMB mentioned, developing sufficient enterprise capacity for AI innovation underscores the critical need for a holistic and robust framework that also entails appropriate funding and policy development capabilities. These resources are needed to

establish a resilient and future-ready AI ecosystem and face AI development, testing, and deployment challenges and opportunities.

2. Removing Barriers to the Responsible Use of Artificial Intelligence

a. **IT Infrastructure-** To ensure that the AI projects have access to adequate IT infrastructure, including satisfactory access to software tools, open-source libraries, and deployment and monitoring capabilities required to rapidly develop, test, and maintain AI applications, technology capabilities, and automation should be introduced for allocating and administering IT resources as computing power, storage, and networking, responding dynamically to the evolving requirements of AI projects. This approach enhances efficiency by simplifying resource allocation, minimizing manual intervention, and guaranteeing timely access to the necessary resources for AI projects.

   Further, establish and develop self-service platforms that provide AI developers with easy access to software tools, deployment, and monitoring capabilities. These platforms can automate tasks like tool installation, configuration, and access management, empowering developers to focus on their work without unnecessary hurdles.

b. **Data-** With the aim to establish frameworks and protocols that securely share data within the agency and with authorized external partners, the OMB should encourage measures for safeguarding sensitive data and upholding principles of data privacy while developing centralized platforms for overseeing data access, usage, and security policies in support of data sharing essential to AI applications. Certain platforms and technology solutions could streamline data governance tasks, ensure compliance with data privacy regulations, enhance data usage transparency, and employ advanced tools to efficiently curate and label extensive datasets for AI training. This streamlined approach expedites data preparation, minimizes the potential for human error, and ensures consistency in data quality adherence to data privacy principles.

c. **Cybersecurity-** Agencies are urged to enhance cybersecurity measures' adoption processes, budgeting, and authorizations, particularly for AI applications, by advancing the adoption of continuous authorizations. In alignment with the AI Executive Order, officials should prioritize security controls for generative AI and other critical AI emerging threat vectors. Relevant oversight procedures should be developed to ensure agency adoption of measures, and such measures should integrate automated tools and processes (including in the context of ensuring AI software development lifecycle (SDLC) security) so agencies can deploy controls for AI applications for vulnerabilities and perform measures such as penetration testing and AI red teaming. This automation can identify potential security flaws early, enabling timely remediation and reducing the risk of cyberattacks.

➢ **Minimum Practices for Either Safety-Impacting or Rights-Impacting AI**

1. **Complete an AI impact assessment-**

a. The OMB Memo outlines essential practices for agencies deploying AI, emphasizing risk assessment, stakeholder consideration, data quality evaluation, real-world performance testing, independent evaluation, ongoing monitoring, and human oversight.

b. As mentioned, Agencies must assess potential risks, document stakeholders impacted by AI, and consider failure modes, especially for underserved communities; they should also implement automated tools and technology contextualization techniques to identify and prioritize potential risks associated with AI applications, analyze AI models, data sets, and intended use cases to identify potential biases, security vulnerabilities, and ethical concerns. Additionally, developing a comprehensive risk assessment framework should consider both technical and non-technical risks related to AI, including factors such as data privacy, algorithmic bias, fairness, explainability, and potential impacts on underserved communities.

c. Agencies are required to thoroughly evaluate the quality and relevance of data used in AI design, considering factors like provenance, relevance to the task, breadth, reliability, and error measurement, and thus implementing continuous risk monitoring mechanisms to track and evaluate the evolving risks associated with AI applications as they are deployed and used, including real-time data analysis, anomaly detection, and user feedback mechanisms. However, more focus should be given here to the use of privacy-enhancing technology, data security, and data provenance measures, as well as advanced methods of encryption and state-of-the-art cryptography.

d. A key focus should be given to establishing transparent and automated risk reporting, testing, and mitigation strategies to communicate identified risks and mitigation strategies to relevant stakeholders, including agency leadership, policy experts, and affected communities, to foster trust and accountability in AI cross-sector participation and ecosystem. Measurement of agency posture should be tied to accountability measures to support more robust adoption of controls.

e. Although continuous monitoring with periodic human reviews, at least annually, is essential to detect AI functionality degradation or changes in impact, agencies should also defend against AI-specific exploits by implementing automated or tech-based failure mode analysis (FMEA) techniques to identify potential failure modes in AI systems and their broader environments, considering human factors, software vulnerabilities, components, red-teaming, external dependencies, and unexpected inputs.

➢ **Managing Risks in Federal Procurement of Artificial Intelligence**

a. The OMB Memo outlines key considerations for federal agencies in procuring AI, emphasizing alignment with measures needed to protect security, enhance transparency, competition promotion, data value maximization, and AI-responsible procurement. Steps should be taken to ensure transparency and performance of procured AI, including obtaining adequate documentation by evaluating performance claims and developing automated interoperability testing frameworks to evaluate the compatibility of AI solutions with existing systems and ensure seamless integration across different vendors and technologies, and considering contracting provisions for continuous improvement. As well as establishing automated data protection monitoring mechanisms and data sharing protocols to track and enforce compliance with data protection regulations and ensure the security and privacy of sensitive data throughout the AI lifecycle. All of this requires a more technology-scaled approach to compliance and measurements.

As the Memo and its implementation evolve, we look forward to discussing these proposals with OMB and are available for any questions. We remain excited to collaborate with OMB to increase engagement with innovative companies. We note that technology tools for supporting at-scale risk management and measurements of controls are required under the Memo. Below, we included additional detailed feedback on the Memo.

Respectively,

Dr. Amit Elazari, CEO & Co-Founder, OpenPolicy