NARA AI Compliance Plan for OMB Memorandum M-24-10 – September 2024

- Prepared by **Gulam Shakir**, Chief Technology Officer (**CTO**) and Chief Artificial Intelligence Officer (**CAIO**)
- Issued by **Colleen Shogan, Archivist of the United States**

1. STRENGTHENING AI GOVERNANCE

General

Describe any planned or current efforts within your agency to update any existing internal AI principles, guidelines, or policy to ensure consistency with M-24-10.

In response to OMB memo M-24-10, we are currently undertaking a comprehensive review of our existing internal Artificial Intelligence (AI) principles, guidelines, and policies.

We established our NARA AI Strategic Goals and Objectives, which were approved by our Senior Leadership/Management Team. Our AI Strategic Goals are:

- Goal 1: Achieve NARA's Business Mission Through AI Integration
- Goal 2: Establish Responsible and Ethical Use of AI at NARA
- Goal 3: Foster a Culture of Innovation at NARA
- Goal 4: Strengthen and Upgrade NARA's Security and Infrastructure to Support the Expanded Use of AI/Machine Learning (ML)

We will establish a cross-functional working group, composed of representatives from information technology (IT), business, legal, security and privacy, records management, and other relevant departments, to lead this effort. This group will analyze M-24-10, identifying areas where our existing frameworks may require updates or enhancements to ensure full compliance. Our focus areas include:

- **Strengthening AI Governance:** We're evaluating our existing governance structures to ensure they adequately address the memo's requirements for accountability, transparency, and risk management.
- **Advancing Responsible AI Innovation:** We're reviewing our AI development and deployment processes to incorporate principles of fairness, equity, and non-discrimination, as emphasized in the memo.
- **Managing AI Risks:** We're enhancing our risk assessment and mitigation procedures, particularly for AI systems that could impact public rights or safety.

We plan to complete this review and implement necessary updates by FY 2025. We also anticipate an ongoing refinement of our AI principles and guidelines as the field of AI continues to evolve and as we gain more experience with its practical applications within NARA. Our goal is to maintain a dynamic and adaptable framework that ensures responsible and ethical AI use while supporting NARA's mission.

AI Governance Bodies

Identify the offices that are represented on your agency's AI governance body. Describe the expected outcomes for the AI governance body and your agency's plan to achieve them.

While the establishment of an AI Governance Body is not mandatory for non-CFO Act agencies, NARA is proactively taking this step to align with AI best practices and ensure responsible and ethical AI development and deployment.

.

The following offices are involved in NARA's AI Governance Board:

- Agency Services: Represents the records management, federal records centers, and national declassification centers where AI is being or will be used
- Business Support Services: Represents the operational areas for facility and property management where AI is being or will be used
- Office of the Chief Information Officer: Oversees IT strategy and implementation, including AI infrastructure and security
- Office of the Chief Technology Officer / Chief AI Officer: Leads the technical aspects of AI development and deployment
- Office of General Counsel: Provides legal guidance on AI-related matters, including privacy, ethics, and compliance
- Office of Innovation: Represents the digital engagement and innovations where AI is being or will be used
- Records Management Office: Ensures AI aligns with archival standards and practices
- Research Services: Represents a key stakeholder in using AI
- Presidential Libraries: Represents a key stakeholder in using AI

NARA's AI governance body is expected to achieve effective AI governance, as well as responsible AI innovation, risk management, and collaboration. NARA plans to achieve this by establishing a clear governance structure, developing and implementing AI policies, conducting risk assessments, promoting transparency, fostering collaboration, and providing training. NARA will:

- formalize its AI governance body, ensuring clear accountability and decision-making authority; and will establish our AI governance body charter
- create and implement comprehensive AI policies and procedures to address data governance, algorithm development, and ethical considerations. These policies will guide the responsible development and use of AI throughout the organization
- conduct appropriate risk assessments for all AI systems, enabling proactive risk identification and mitigation
- implement mechanisms to ensure transparency and accountability in AI development and deployment, including documentation, explanations, and feedback channels
- encourage collaboration and communication among stakeholders involved in AI initiatives to facilitate knowledge sharing and coordinated decision-making

- provide ongoing training and education to staff on AI-related topics to ensure responsible and effective AI use

Describe how, if at all, your agency's AI Governance body plans to consult with external experts as appropriate and consistent with applicable law. External experts are characterized as individuals outside your agency, which may include individuals from other agencies, federally funded research and development centers, academic institutions, think tanks, industry, civil society, or labor unions.

NARA's AI governance body is committed to responsible AI use and plans to consult external experts when necessary and legally permissible. This includes engaging subject matter experts, soliciting public feedback, and collaborating with other agencies. By leveraging external perspectives, NARA aims to ensure its AI initiatives align with best practices, address potential challenges, and meet societal expectations.

Potential avenues for external consultation include:

- **Engaging subject matter experts:** NARA will seek guidance from recognized experts in AI ethics, law, policy, and technical implementation. This could involve collaborating with academic institutions, or participating in industry forums and workshops.
- **Participating in public engagement:** NARA will actively solicit feedback from the public and relevant stakeholders on AI initiatives, ensuring transparency and incorporating diverse perspectives into decision-making processes.
- **Collaborating with federal agencies and other entities:** NARA will collaborate with other federal agencies, research communities, and academic institutions to share best practices on AI governance and implementation, fostering a community of learning and innovation.

Describe your agency's process for soliciting and collecting AI use cases across all sub agencies, components, or bureaus for the inventory. In particular, address how your agency plans to ensure your inventory is comprehensive, complete, and encompasses updates to existing use cases.

The Office of the CIO spearheaded the collection of AI use cases by collaborating with key stakeholders across several business areas. We will implement a centralized intake form to collect AI use cases across NARA, conduct targeted outreach to stakeholders, foster cross-functional collaboration, and iteratively refine the process. This approach will ensure comprehensive representation and generate a valuable inventory for future AI initiatives. Additionally, we will create a Standard Operating Procedure (SOP) to maintain and update the AI use cases inventory.

By leveraging the expertise of external stakeholders, we're confident that our AI governance body can make informed decisions, promote responsible AI use, and ensure that NARA's AI initiatives serve the public good.

Describe your agency's process for soliciting and collecting AI use cases that meet the criteria for exclusion from being individually inventoried, as required by Section 3(a)(v) of M-24-10. *In particular, explain the process by which your agency determines whether a use case should be excluded from being individually inventoried and the criteria involved for such a determination.*

At NARA, we're committed to responsible AI adoption in line with OMB M-24-10. To identify AI use cases eligible for exclusion from the inventory, we employ a systematic process. We broadly

solicit use cases from all relevant agency units, then rigorously assess them against the exclusion criteria.

For those meeting the criteria, we maintain detailed documentation and conduct expert reviews to ensure accuracy and consistency. Finally, we implement continuous monitoring to address any changes in projects that might affect their exclusion status. This process helps us maintain transparency while safeguarding sensitive information and internal processes, demonstrating our dedication to responsible AI governance.

Identify how your agency plans to periodically revisit and validate these use cases. ***In particular, describe the criteria that your agency intends to use to determine whether an AI use case that previously met the exclusion criteria for individual inventorying should subsequently be added to the agency's public inventory.***

NARA will periodically revisit and validate excluded AI use cases to ensure they remain compliant with M-24-10. This involves regular reviews of documentation and expert evaluations. Stakeholder input will be gathered, and the AI governance body will make final determinations on exclusion status.

2. ADVANCING RESPONSIBLE AI INNOVATION

NARA will invest in AI research and development to advance responsible AI technologies. NARA will establish an AI Ethics Review Team to review and assess the ethical implications of AI projects. We will develop clear AI Ethics Guidelines to inform AI development and deployment. We will collaborate with external stakeholders to share best practices and promote responsible AI innovation. NARA will conduct periodic audits and assessments of AI systems to ensure compliance.

Removing Barriers to the Responsible Use of AI

Describe any barriers to the responsible use of AI that your agency has identified, as well as any steps your agency has taken (or plans to take) to mitigate or remove these identified barriers. ***In particular, elaborate on whether your agency is addressing access to the necessary software tools, open-source libraries, and deployment and monitoring capabilities to rapidly develop, test, and maintain AI applications.***

NARA recognizes potential barriers to responsible AI use, including data quality and bias, transparency, privacy and security, and skills gaps. To address these, we implement robust data governance, explainable AI, privacy by design principles, workforce development, public engagement, and ongoing evaluation. By proactively mitigating risks, NARA will ensure AI is used ethically and effectively.

Identify whether your agency has developed (or is in the process of developing) internal guidance for the use of generative AI. ***In particular, elaborate on how your agency has established adequate safeguards and oversight mechanisms that allow generative AI to be used in the agency without posing undue risk.***

NARA will establish a cross-functional working group to research and understand generative AI. Based on these insights, NARA will develop comprehensive guidelines outlining permissible and prohibited use cases, while also providing training to ensure staff understand the guidelines and potential risks. NARA will continually refine the guidance, ensuring it stays current with the evolving AI landscape, and will maintain transparency by making the guidance easily accessible

and communicating updates regularly. This systematic approach aims to create a framework that empowers NARA staff to leverage generative AI responsibly and ethically.

AI Talent

Describe any planned or in-progress initiatives from your agency to increase AI talent. *In particular, reference any hiring authorities that your agency is leveraging, describe any AI focused teams that your agency is establishing or expanding, and identify the skillsets or skill levels that your agency is looking to attract. If your agency has designated an AI Talent Lead, identify which office they are assigned to.*

NARA is actively enhancing its AI talent through the targeted recruitment of experts, upskilling existing staff via training programs, and fostering internal AI communities. We're also creating hands-on AI experiences through rotational programs and collaborations, ultimately cultivating a workforce capable of leveraging AI's potential to drive innovation and fulfill NARA's mission. NARA has created position descriptions (PDs) for an Artificial Intelligence Specialist and Senior Artificial Intelligence Specialist.

If applicable, describe your agency's plans to provide any resources or training to develop AI talent internally and increase AI training opportunities for Federal employees. *In particular, reference any role-based AI training tracks that your agency is interested in, or actively working to develop (e.g., focusing on leadership, acquisition workforce, hiring teams, software engineers, administrative personnel or others).*

NARA prioritizes building AI talent through targeted training, mentorship, and hands-on projects. We will partner with educational institutions and industry leaders to offer specialized courses and certifications, while also encouraging participation in federal AI training programs. This multi-pronged approach will create a skilled workforce capable of driving innovation and fulfilling NARA's mission in the digital age.

AI Sharing and Collaboration

Describe your agency's process for ensuring that custom-developed AI code—including models and model weights—for AI applications in active use is shared consistent with Section 4(d) of M-24-10.

NARA will maintain an inventory of custom-developed AI code and applications, assessing their impact and classifying them accordingly. Any safety or rights impacting applications will be carefully evaluated before sharing their code, models, and weights, ensuring compliance with laws and policies. Non-high-impact applications are generally shared unless restrictions apply.

Currently, we're not developing any custom models; rather, we're leveraging the existing large language models (LLMs) and developing code to access those LLMs. We will use secure mechanisms for sharing and maintaining transparent documentation of our decisions by establishing a SOP. This process will be periodically reviewed for consistency with M-24-10.

Elaborate on your agency's efforts to encourage or incentivize the sharing of code, models, and data with the public. Include a description of the relevant offices that are responsible for coordinating this work.

NARA actively encourages the use of open-source software and tools, fostering transparency and community collaboration. NARA will consistently evaluate open-source tools to determine if they meet our requirements. We will leverage secure data sharing platforms to make anonymized and/or de-identified datasets available to the public, prioritizing privacy protection. We will document our internal process/workflow for public data and code sharing involving

expertise from NARA's Information Services, Innovation, Security, and Privacy teams. Through public engagement, we will further promote the use of our shared AI resources.

Harmonization of Artificial Intelligence Requirements

Explain any steps your agency has taken to document and share best practices regarding AI governance, innovation, or risk management. ***Identify how these resources are shared and maintained across the agency.***

NARA is considering several initiatives to document and share best practices regarding AI governance, innovation, and risk management:

- **Internal Knowledge Sharing:** We've established internal communities of practice and knowledge repositories to facilitate the sharing of lessons learned, successful use cases, and emerging trends in AI governance and innovation.
- **AI Use Case Inventory:** The development and publication of the AI Use Case Inventory on Archives.gov and AI.gov serves as a valuable resource, not just for NARA, but also for other agencies and the public to learn about potential applications of AI in the archival and cultural heritage sector.
- **Collaboration with Federal Partners:** We actively engage with other federal agencies through interagency working groups and forums to exchange insights and best practices in AI governance, innovation, and risk management.
- **External Engagement:** We participate in relevant conferences, workshops, and webinars to share our experiences and learn from others in the field. We also publish articles and blog posts on our website to disseminate information on our AI initiatives and best practices.

We're committed to continuously refining our approach and expanding our outreach efforts to ensure that our knowledge and experience benefit the broader AI community.

3. MANAGING RISKS FROM THE USE OF ARTIFICIAL INTELLIGENCE

Determining Which Artificial Intelligence Is Presumed to Be Safety-Impacting or Rights-Impacting

Explain the process by which your agency determines which AI use cases are rights-impacting or safety-impacting. ***In particular, describe how your agency is reviewing or planning to review each current and planned use of AI to assess whether it matches the definition of safety-impacting AI or rights-impacting AI, as defined in Section 6 of M-24-10. Identify whether your agency has created additional criteria for when an AI use is safety-impacting or rights-impacting and describe such supplementary criteria.***

NARA will establish a cross-functional team comprising experts in AI, ethics, cybersecurity, privacy, law, and relevant subject matter areas that will be responsible for assessing AI use cases. The cross-functional team will review each AI use case in the inventory and assess whether it meets the definitions of "safety-impacting" or "rights-impacting" as defined in Section 6 of M-24-10.

At NARA, we do not have safety-impacting or rights-impacting AI use cases. Because NARA is not a law enforcement, public safety or benefits granting agency, we do not

anticipate having these use cases. However, we will monitor proposed use cases of AI and implement the required protections and steps should a use-case for such AI emerge.

Implementation of Risk Management Practices and Termination of Non-Compliant AI

Describe your agency's process for issuing, denying, revoking, tracking, and certifying waivers for one or more of the minimum risk management practices.

Elaborate on the controls your agency has put in place to prevent non-compliant safety impacting or rights-impacting AI from being deployed to the public.

For major information systems, NARA conducts security assessments of any applications it develops or leverages that it makes available to the public, including AI applications. These assessment processes include scanning custom developed code with Static Application Security Testing (SAST) tool, scanning front end applications with a Dynamic Application Security Testing (DAST) tool, and manual review and testing.

Minimum Risk Management Practices

Identify how your agency plans to document and validate implementation of the minimum risk management practices. ***In addition, discuss how your agency assigns responsibility for the implementation and oversight of these requirements.***

Where required, NARA will document their AI risk management efforts through clear procedures, detailed risk registers, and comprehensive audit reports. The validation of these practices will be achieved through technical reviews, cybersecurity reviews and continuous monitoring of AI systems, and gathering feedback from users and stakeholders.

Responsibility for the successful implementation and oversight will be shared across the agency: AI system owners will take the lead in implementing risk management practices, risk management teams will provide guidance and oversight, senior leadership will maintain ultimate accountability, and cross-functional collaboration will be essential to address risks comprehensively.