Social Security Administration (SSA) Compliance Plan for OMB Memoranda M-24-10 – September 2024

Prepared by Brian Peltier, Chief AI Officer (CAIO)

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.

The Social Security Administration (SSA) has taken proactive steps to update its internal AI principles, guidelines, and policies to align with the Office of Management and Budget Memo M-24-10 (OMB M-24-10). The Chief Information Officer (CIO) has issued policy emphasizing the use of responsible AI throughout the agency. This policy establishes enterprise-wide principles and mandates compliance with OMB guidance.

To ensure responsible AI practices, SSA has formed the Responsible AI Core Team (RAI Core Team). This team oversees the implementation of the Responsible AI Implementation Framework (RAI Framework), which assists teams in evaluating AI use cases for compliance with OMB guidance and promotes best practices related to all the principles of responsible AI usage across all stages of the AI lifecycle.

Furthermore, the RAI Core Team is introducing the AI Intake process, which enables SSA to stay informed about emerging AI use cases under consideration. This process facilitates effective monitoring and evaluation of nascent AI initiatives within the agency.

These initiatives collectively aim to foster the responsible and trustworthy use of AI within SSA, while upholding compliance with relevant regulations and promoting transparency in AI practices.

AI Governance Bodies

- Identify the offices that are represented on your agency's AI governance body.

The SSA AI Senior Executive Council (AI SEC) Charter establishes membership of the agency's AI governance body as follows:

  - Office of the General Counsel (OGC)
  - Office of Transformation (OT)
  - Office of Analytics, Review, and Oversight (DCARO)
  - Office of Budget, Finance, and Management (DCBFM)
  - Office of Civil Rights and Equal Opportunity (DCCREO)
  - Office of Hearings Operations (DCHO)
  - Office of Human Resources (DCHR)
  - Office of Operations (DCO)

  - Office of Retirement and Disability Policy (DCRDP)
  - Office of the Chief Information Officer (OCIO) (including the CAIO)

- Describe the expected outcomes for the AI governance body and your agency's plan to achieve them.

The AI SEC, consisting of the SSA Deputy Commissioners and other senior leadership from the mentioned offices, plays a crucial role in providing executive-level guidance for managing AI risk and facilitating the implementation of AI initiatives.

The AI SEC's primary objective is to ensure the responsible use of AI within the agency. This involves leveraging AI technologies to enhance operational efficiencies while upholding ethical and legal standards. The AI SEC creates an environment that allows AI use case teams to operate with flexibility, enabling SSA's limited AI resources and expertise to focus on AI innovation and risk management.

Furthermore, the AI SEC ensures that various aspects such as IT infrastructure, data, cybersecurity, and legal, privacy, and ethical considerations are thoroughly addressed across the enterprise. This includes implementing appropriate institutional safeguards to protect the integrity and security of AI-related activities.

- 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.

The AI SEC recognizes the importance of ensuring comprehensive and informed decision-making in AI governance. To facilitate this, the AI SEC and RAI Core Team engage in regular consultation with relevant experts. This collaborative approach ensures that SSA benefits from the expertise and insights of external experts while developing and implementing AI projects.

Additionally, SSA actively attends and presents at conferences related to AI (such as this year's presentations to the National Disability Forum regarding how AI may affect the landscape of Social Security), providing transparency and opportunities to engage with external experts and exchange knowledge and experiences to the extent permitted by applicable law.

By actively seeking input and collaboration, the AI governance body within the agency ensures that its decision-making processes are well-informed, consistent with applicable law, and benefit from a diverse range of perspectives.

AI Use Case Inventories

- 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*.

To ensure a comprehensive and complete inventory of AI use cases, the RAI Core Team, acting on behalf of the CAIO, utilizes agency standard communication channels to reach out to every deputy commissioner's office and sub-component at SSA to collect input across the agency on AI use cases.

This approach ensures that all levels of management throughout the agency are well-informed about AI governance, the inventory, can respond to it, and are accountable for maintaining its accuracy.

Furthermore, the RAI Core Team follows up with additional questions to ensure that the responses provided are comprehensive and meet the intended purpose of the inventory. This iterative process helps capture any updates or changes to existing use cases, while ensuring that the inventory remains up-to-date and reflective of the agency's AI initiatives.

Once the responses are collected, the RAI Core Team compiles the results and presents them to the CAIO and the AI SEC for review and concurrence prior to publication. This review process ensures that the inventory is accurate, comprehensive, complete, and provides a reliable snapshot of the agency's AI use cases.

Reporting on AI Use Cases Not Subject to Inventory

- 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.*

The process for soliciting and collecting AI use cases that meet the criteria for exclusion from individual inventory reporting involves collaboration between the CAIO, the RAI Core Team, and the use case owners.

The CAIO, in consultation with the RAI Core Team, conducts evaluations to determine whether an AI use case should be excluded from being individually inventoried based on OMB instructions (such as when the details of being individually inventoried could potentially compromise its effectiveness, such as in the case of security or anti-fraud solutions). This evaluation considers OMB criteria to ensure that the exclusion is appropriate and justified, while still including aggregate metrics about the use case in the OMB AI Inventory.

It is important to note that even if a use case is excluded from individual inventory reporting, it is still actively monitored and tracked. This ensures that the agency maintains awareness of all AI initiatives, including those that are not individually inventoried, and can effectively manage and assess their impact, compliance with applicable requirements, and risk posture.

By considering the relevant criteria, the agency can accurately determine if AI use cases should be excluded from individual inventory reporting while still maintaining oversight and monitoring of these use cases.

- 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.*

As part of the annual AI inventory process, SSA actively monitors any use cases that were previously excluded from individual inventorying. This monitoring includes updating information related to these use cases for reevaluation. By revisiting these use cases on a regular basis, the agency can assess whether any changes have occurred that warrant their inclusion in the public inventory.

Additionally, the RAI Core Team maintains regular communication with the owners of excluded use cases. This communication serves as a means to stay informed about any substantive changes that may have taken place. If a significant change occurs in an excluded use case that could impact its eligibility for exclusion, the RAI Core Team can initiate a reevaluation process.

By periodically revisiting and validating excluded use cases, SSA can maintain an accurate and up-to-date public inventory that reflects the agency's AI initiatives and provides transparency to stakeholders.

2. ADVANCING RESPONSIBLE AI INNOVATION

AI Strategy [Optional]

- [OMB context note: Please note that the agency AI Strategy requirement for CFO Act agencies identified in Section 4(a) of M-24-10 is separate and distinct from this compliance plan. Reporting on the AI Strategy is not required in this compliance plan, but agencies can include information on the forthcoming strategies if desired.]

SSA will provide the AI Strategy as identified in OMB M-24-10 by March 28, 2025.

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.*

SSA has identified certain barriers to the responsible use of AI, including the need for IT architectural patterns. To accelerate the removal of these barriers, the development of standard architectural patterns is underway for AI technology to provide AI use cases with a structured approach to designing and organizing necessary IT infrastructure, ensuring that AI technology solutions align with the organization's goals, security requirements, and legal and regulatory compliance, which includes the responsible use of AI.

In addition, the agency has developed the RAI Framework, which coincides with the architectural patterns, to provide guidance to AI investigation teams and AI project teams. The RAI Framework assists these teams in managing and addressing potential risks associated with AI initiatives at all stages of the development lifecycle. By following the guidance provided in the RAI Framework, SSA aims to promote responsible and ethical AI practices while mitigating any potential barriers to the responsible use of AI.

Finally, SSA recognizes that awareness, knowledge, and trust of AI may cause a barrier to the responsible use of AI. The RAI Core Team has created trainings, presentations, and documentation, to supplement the RAI Framework and to guide various stakeholders in developing understanding as to their role in responsible AI at SSA. Specifically, the RAI Core Team works with executives to understand their role in driving agency-wide adoption of Responsible AI principles. The RAI Core Team works with use case teams to help them understand their responsibilities related to the RAI Framework and EOs. Finally, the RAI Core Team works with line employees to help them understand the benefits and drawbacks of AI and how to use AI safely and responsibly.

- 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.*

Due to the potential severe consequences that could arise from unauthorized data disclosure, protecting SSA's data assets is of utmost importance. As a data-centric organization, SSA places a high priority on the security, integrity, and privacy of its systems and data. The consequences of data leakage, both into and out of the agency, highlight the critical need for safeguarding data assets.

To mitigate the risks associated with generative AI and protect SSA data, SSA has implemented several safeguards and oversight mechanisms. The agency has blocked general employee access to external third-party generative AI services, recognizing that such access could result in the unauthorized disclosure of SSA data, including Personally Identifiable Information. This decision ensures that adequate safeguards are in place to protect SSA data from being exposed to or further used by third-party generative AI systems.

In addition, the RAI Core Team has developed and distributed optional training on the risks associated with generative AI. This training helps inform employees within the agency about potential risks and the need for caution when utilizing generative AI technologies. By raising awareness and providing guidance, SSA aims to ensure that employees understand the risks and take appropriate measures to protect data assets when using generative AI.

SSA acknowledges the opportunities presented by AI, including increased efficiency, higher quality, more equitable service, and more meaningful work for staff. The agency is committed to responsibly leveraging cutting-edge technologies to enhance operations and improve public service. SSA will continue conducting research to identify suitable use cases for AI and evaluate the effects of its implementation.

By implementing these safeguards, oversight mechanisms, and comprehensive training, SSA aims to strike a balance between leveraging the benefits of generative AI and protecting the security and privacy of its data assets. The agency remains committed to ensuring that generative AI is used responsibly and without posing undue risk.

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.

SSA is leveraging Direct Hire Authority to recruit data scientists. The SSA office leading the RAI Core Team has been allocated 3 new positions enabling the agency to attract individuals with specialized AI expertise.

Furthermore, SSA is actively working on repositioning employees to develop AI talent internally. This approach allows existing employees to acquire the necessary skills and knowledge in AI, fostering a culture of continuous learning and growth within the agency.

To support the development of AI talent, SSA offers a range of AI trainings through its internal training system. These trainings are designed to provide employees with role-based AI knowledge and skills, ensuring they have the necessary competencies to contribute effectively to AI initiatives.

In addition to training opportunities, SSA maintains an active AI Community of Interest. This community serves as a platform for employees to share information, collaborate, and foster the growth of AI talent within the agency. It provides a space for knowledge exchange and encourages employees interested in AI to connect and learn from one another.

- 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).*

See previous response.

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.

In line with its commitment to transparency and collaboration, SSA will review each use case to determine its eligibility to share source code publicly to the extent permitted by applicable law through the SSA Open Source Code Site and the Developer Support Site, in accordance with the SSA Open Government Plan.

- 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.

SSA upholds its commitment to transparency and open government by making code publicly available to the extent permitted by applicable law through the SSA Open Source Code Site and the Developer Support Site. This practice aligns with the SSA Open Government Plan.

Similarly, the Chief Data Officer, in coordination with other agency senior officials, ensures the public sharing of data through the Social Security Data Page, in accordance with the M-13-13 Open Data Policy. These initiatives promote accessibility and facilitate the use of SSA's code and data by the public, to the extent permitted by applicable law.

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.*

SSA has implemented various steps to document and share best practices regarding AI governance, innovation, and risk management throughout the agency.

The RAI Core Team plays a vital role in this process by maintaining the RAI Framework and the AI Community of Interest. The AI Community of Interest serves as a valuable resource that contains best practices, blog posts, and updates on AI governance. This community ensures that employees have access to updated information and guidance related to AI. It fosters collaboration and knowledge exchange among employees interested in AI, facilitating the sharing of best practices and lessons learned. The RAI Framework, on the other hand, provides a comprehensive set of requirements and considerations for AI use cases, promoting consistency and responsible AI implementation across the agency.

The CAIO represents SSA in the CAIO Council, actively sharing information and experiences with other federal government entities. This participation allows SSA to contribute to the broader AI community, fostering collaboration and knowledge exchange with other agencies. The CAIO and RAI Core Team members also actively participate in the working groups of the CAIO Council, further contributing to the sharing of best practices and experiences.

Additionally, the Chief Architect plays a crucial role in establishing and maintaining enterprise architecture patterns specifically tailored for AI implementations. These patterns ensure consistency and efficiency across the agency's AI initiatives, providing a framework for effective AI governance and risk management.

By leveraging these resources, platforms, and participation in external forums, SSA ensures that best practices regarding AI governance, innovation, and risk management are documented, shared, and maintained across the agency. This promotes consistency, collaboration, and continuous improvement in the agency's AI initiatives, fostering responsible and effective use of AI technologies.

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.*

The RAI Core Team gathers necessary information and works closely with use case owners to assess each AI use case. The CAIO, in collaboration with the RAI Core Team, conducts evaluations to make determinations on whether AI use cases are rights-impacting or safety-impacting. These evaluations are conducted in accordance with federal mandates, including EO 14110 and OMB M-24-10. The agency ensures that internal and external experts, as well as use case owners, provide input during the determination process.

For rights-impacting determination, the agency evaluates use cases based on whether the AI output serves as the "principal basis for a decision or action." Next, the agency reviews whether the decision or action involves one or more specific individuals or entities. Finally, the agency reviews whether the decision or action has a "legal, material, binding, or significantly significant effect on an individual or entity's civil rights, civil liberties, privacy, equal opportunities, or access to critical government resources or services." This ensures that the agency assesses whether the AI use case is rights-impacting, consistent with the definitions of rights-impacting AI as defined within M-24-10.

Regarding safety-impacting AI, the agency is continually evaluating its use cases to evaluate where potentially safety-impacting AI could occur consistent with M-24-10 and will make those determinations and document them when necessary.

- If your agency has developed its own distinct criteria to guide a decision to waive one or more of the minimum risk management practices for a particular use case, describe the criteria

SSA has not developed its own criteria. SSA follows the criteria outlined in EO 14110 and OMB M-24-10.

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

The CAIO, in collaboration with the RAI Core Team, will conduct evaluations and make determinations based on the requirements set forth in EO 14110 and OMB M-24-10. These evaluations will consider the specific circumstances and needs of each system or use case.

Based on these evaluations, the CAIO, supported by the RAI Core Team, will determine whether to issue, deny, or certify waivers for each use case's minimum risk management practices. This decision-making process ensures that appropriate considerations are made to balance the need for flexibility with the importance of risk management.

Any waivers that are granted will be properly tracked and documented in the AI inventory. This tracking mechanism ensures that waivers are accounted for and can be monitored for compliance and ongoing evaluation.

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

- 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.

SSA has implemented several controls to prevent the deployment of any non-compliant AI, including safety-impacting or rights-impacting AI, to the public. These controls are designed to ensure adherence to the guidelines and principles outlined in EO 14110 and OMB M-24-10.

The RAI Framework, which is the foundation of SSA's AI governance, includes risk management activities and required risk artifacts. These activities help identify, assess, and mitigate potential risks associated with AI initiatives. By incorporating compliance checkpoints within the framework, SSA ensures that safety-impacting or rights-impacting AI applications undergo thorough evaluation and review before being deployed to the public.

Furthermore, SSA adheres to Authority to Operate (ATO) processes. These processes ensure that only IT solutions that meet stringent criteria for safety, security, privacy, legality, and alignment with national values are authorized for deployment to the public. The ATO processes provide an additional layer of control and oversight to prevent the deployment of non-compliant AI applications.

- Describe your agency's intended process to terminate, and effectuate that termination of, any non-compliant AI.

The CAIO, in collaboration with relevant stakeholders, will assess the impact of the termination and determine the appropriate time frame for the removal.

During the assessment, the CAIO will consider the potential impact of the termination on vital processes and operations. The goal is to ensure that the termination is carried out in a manner that minimizes disruption and any negative consequences.

Once the decision to terminate a non-compliant AI application is made, the necessary steps will be taken to effectuate the termination. This may involve disabling or removing the AI application from the systems and ensuring that any associated data or dependencies are appropriately handled.

SSA aims to terminate any non-compliant AI applications, ensuring compliance promptly and effectively with relevant guidelines, and minimizing any potential negative impact on the agency's operations.

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.

SSA has established a process to document and validate the implementation of the minimum risk management practices for all AI use cases. The responsibility for implementation and oversight of these requirements is assigned to the AI use case owners.

AI use case owners are responsible for identifying, assessing, and mitigating risks that are specific to their individual use cases. They follow established enterprise risk management practices as discussed further below to ensure a comprehensive approach to risk management. By adhering to these practices, they can effectively identify potential risks associated with AI implementation and take appropriate measures to mitigate them.

The RAI Framework provides guidance to AI use case owners, assisting them in identifying and mitigating risks that are specific to AI. The RAI Framework requires use cases to create governance artifacts that document their compliance with the processes outlined in the RAI Framework. The RAI Framework also requires the completion of a Bias Mitigation Report for rights and safety impacting use cases.

The RAI Core Team will play an active role in monitoring the risk management practices of AI use cases. It will ensure that the use case owners are complying with the minimum risk management practices and that these practices are effective in mitigating risks associated with responsible AI principles.

Furthermore, the CAIO and AI SEC will have the responsibility to oversee and monitor systemic risks associated with the use of AI across the agency. They will ensure that appropriate measures are in place to address and mitigate these risks at an agency-wide level. This oversight helps to ensure that risk management practices are consistently implemented and that potential risks are identified and addressed in a timely manner.

By sharing responsibility with AI use case owners and implementing oversight mechanisms, SSA ensures that the comprehensive risk management practices for AI use cases are documented, validated, and effectively implemented throughout the agency.