



Board of Governors of the Federal Reserve System Compliance Plan for OMB Memorandum M-24-10

September 2024



Overview

In accordance with the Office of Management and Budget's (OMB) Memorandum M-24-10, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence, the Board of Governors of the Federal Reserve System (Board) is pleased to share its plan for compliance with the requirements of M-24-10.

The Board is committed to an artificial intelligence (AI) program for its staff that fosters responsible AI innovation, maintains robust AI governance, and manages the risks associated with the use of AI. As part of the Board's enterprise technology and data governance frameworks, the Chief Artificial Intelligence Officer (CAIO) plays an important role in advancing the Board's AI objectives and coordinating the Board's use of AI. The Board anticipates continued growth in AI capabilities as its AI program matures.

This document outlines the Board's compliance plans in satisfaction of the requirements of section 3(a)(iii) of M-24-10 and section 104(c) of the AI in Government Act. The Board will report its individual use-case-specific practices in accordance with sections 5(c)(iv) and (v) of M-24-10 separately through the annual AI use case inventory.

Strengthening AI Governance

As part of the Information Resources Management (IRM) Strategic Plan, under Objective 1.3, "Artificial Intelligence and Machine Learning," the Board previously set a priority for assessing practical applications of traditional AI, machine learning, and generative AI for enhancing data capabilities and increasing operational efficiencies.¹ The keystone for the Board's AI governance is its AI policy, which establishes guardrails for permissible uses of AI (in Board work and Board-delegated functions executed by the Federal Reserve Banks), goals for adoption, risk-management requirements, and governance processes. The Board will update the policy as the AI Program's understanding of potential AI use cases advances and as AI continues to evolve.

¹ See <https://www.federalreserve.gov/publications/files/irm-strategic-plan-20231208.pdf>.

AI Governance Bodies

In early 2024, the Board established an AI Program under the Office of the Chief Operating Officer and appointed a CAIO to lead the program. The AI Program consists of two governing bodies, the AI Program Team and the AI Enablement Working Group, which are critical components of the Board's commitment to ensuring the responsible and ethical use of AI technologies in Board work. Together, they oversee the implementation of AI best practices, technical enablement, employee education, operation of AI systems as well as help ensure compliance with relevant laws, regulations, and internal policies.

The AI Program team, which includes a subset of senior staff along with supporting staff, has responsibility for administering the Board's internal AI policy, including the tracking of all generative AI permissible use cases, coordinating intake of requests for policy exceptions, and reviewing and routing use case exceptions for risk and impact assessments for approval. The AI Enablement Working Group is composed of representatives from the various Board divisions and key compliance offices, including

- Division of Board Members
- Division of Consumer and Community Affairs
- Division of Financial Management
- Division of Financial Stability
- Division of International Finance
- Division of Information Technology
- Legal Division
- Division of Monetary Affairs
- Division of Management
- Office of the Secretary
- Division of Research and Statistics
- Division of Reserve Bank Operations and Payment Systems
- Division of Supervision and Regulation
- Office of the Chief Data Officer
- Office of Cybersecurity and Privacy
- Enterprise Risk Management Program
- Office of Procurement
- Federal Reserve System Liaisons

In addition, enterprise AI investment decisions are reviewed through the Board's Technology Oversight Committee, which is co-chaired by the Chief Operating Officer with executive members from multiple divisions as well as the chief information officer, chief financial officer, chief data officer, and CAIO.

Expected Outcomes

The AI Program works to achieve the following outcomes:

- **Ethical AI deployment.** Ensure all AI systems are developed and deployed in a manner consistent with scientific, legal, and ethical norms.
- **Risk mitigation.** Identify and mitigate potential enterprise risks associated with AI, including biases, unfair outcomes, privacy lapses, data protection issues, copyright infringement, and other harms.
- **Transparency and accountability.** Maintain transparency in AI operations and hold stakeholders accountable for their roles in AI governance.
- **Continuous improvement.** Foster a culture of continuous improvement in AI governance and adoption practices, keeping pace with technological advancements and emerging best practices.
- **Technical enablement.** Evaluate AI technologies, establish enterprise capabilities, support the use of AI to leverage the Board's data as a strategic asset, implement technical guardrails, and ensure information security to advance AI innovation within the institution.
- **Talent development.** Ensure that the workforce has the tools and education to employ AI technologies responsibly and comply with all relevant AI-related policies.

Consultation with External Experts

The CAIO, along with the AI Program including the AI Enablement Working Group, consults with external experts as appropriate and consistent with applicable laws to enhance the robustness of the Board's AI Program and to incorporate diverse perspectives. These consultations may include

- **Academic communities.** Collaborating with researchers and experts from universities and research institutions.
- **Industry leaders.** Engaging with industry experts and vendors to gain insights into AI technologies and services that could be considered for potential procurement.
- **Civil society organizations.** Consulting with non-governmental organizations and other civil society organizations to better understand the potential societal impact of AI.
- **Interagency collaborations.** Coordinating with other federal agencies, particularly other federal financial regulators, to share knowledge and align on best practices for AI governance and operations.

- **International engagement.** Sharing information on AI developments and best practices with other central banks and international organizations in established international forums.

AI Use Case Inventories

The creation and maintenance of AI use case inventories are essential to ensuring that the Board has a comprehensive understanding of how AI technologies are utilized across the organization. This inventory process allows the CAIO to manage AI deployments effectively, identify opportunities for enterprise-level innovation, ensure alignment with the AI policy, mitigate enterprise risks, uphold ethical standards, and meet regulatory requirements.

Process for Soliciting and Collecting AI Use Cases

For the annual AI Use Case Inventory, the Board will conduct the following process:

- The AI Program will lead a proactive outreach initiative to support a comprehensive collection of all AI use cases. All AI users of AI in connection with Board work will document their use cases and submit them to the AI Program.
- The AI Program will review each use case to flag any potential impact on rights and safety and evaluate exclusion from public reporting, as defined by OMB's *Guidance for Creating Agency Inventories of AI Use Cases EO14110*. The AI Program will gather the additional detailed information required for submission of covered use cases to the formal inventory.
- Use cases will be stored in a common repository that allows reporting, storage, and ongoing tracking as part of the process. This will allow current and future AI users, as well as the AI Program team, to periodically revisit, update, and evaluate use cases as they evolve through the product lifecycle until eventual retirement.

Ensuring Comprehensive and Complete Inventory

The Board will maintain the AI Use Case Inventory and will provide periodic updates to AI use cases and project lifecycle status to ensure accurate and transparent reporting in compliance with OMB reporting requirements. Any potential uses of AI will also be made transparent via existing procurement and technology project review processes. Current and future AI users, as well as the AI Program team, will add new use cases, and revisit and update existing use cases as needed. The Board will also conduct periodic validations of the AI Use Case Inventory in alignment with OMB reporting requirements.

Reporting on AI Use Cases Not Subject to Inventory

The Board will collect information on AI use cases at the individual detail level. This process is described in "[Process for Soliciting and Collecting AI Use Cases](#)." Each use case will be evaluated by the AI Program to determine whether it is a covered case and not subject to exclusion based on

the OMB-defined exclusion criteria. As exclusions arise, the Board will prepare aggregate metrics for excluded use cases, periodically review those cases, and validate the use cases as part of its compliance with OMB reporting requirements.

Advancing Responsible AI Innovation

Embracing innovation requires removing unnecessary and unhelpful barriers to the use of AI while retaining the guardrails that ensure its responsible use. The Board is committed to supporting a culture in which innovation can thrive while also managing the risks that come with AI.

Removing Barriers to Responsible Use of AI

The following provides a summary of identified barriers and the steps the Board has taken, or plans to take, to mitigate or remove the barriers.

- **IT Infrastructure and Data**

- **Access to AI compute environments.** Allow staff access to internal computing environments equipped to run AI models and internally publish a list of available, approved AI computing environments. Evaluate on a continual basis emerging alternatives for secure AI computing based on new use cases.
- **Acquisition and development of secure enterprise AI tools/platforms.** Investigate secure, commercially available AI tools and platforms to accelerate AI maturity and adoption for the most common and low risk use cases throughout the enterprise.
- **Access to open-source AI models and libraries.** Allow access to internal open-source AI foundation models and libraries with vetted, acceptable license terms, maintain a list of approved and available models and libraries, and evaluate requests for new models on an ongoing basis.
- **Established paths to production.** Provide and maintain pathways to transition AI uses from pilot to production, establishing appropriate processes for validation, code reviews, and security risk assessments.
- **Data.** Evaluate whether data are permitted for use with AI and seek modifications as needed. Build internal databases of frequently used unstructured data to promote efficiency.

- **Policy and Operations**

- **Ensure compliance with internal AI policy.** Ensure that Board staff are following the Board's internal AI policy, using only approved tools and data, and reporting uses to the Board AI Program. Implement administrative and technical controls to support safer ongoing operations and responsible AI innovation.

- **Monitor AI policy and technology advancements.** Monitor emerging AI trends, regularly evaluate the AI policy for possible adaptation, and seek feedback continuously about desired AI capabilities from across the user population.
- **Ensure value of AI investments.** Evaluate and determine investment levels required to successfully manage AI, such as technology, governance, operations and skill development, and track the return on those investments.
- **Prioritize role-based AI education and training.** Ensure that Board staff have the requisite AI education and training based on their role to increase adoption and awareness of AI technology and platforms.
- **Cybersecurity and Privacy**
 - **Ensure data transparency and privacy safeguards.** Ensure that technology supports and sustains privacy protections, civil liberties, and civil rights when using safety- and rights-impacting AI.
 - **Incorporate new AI standards, security controls, and guardrails.** Update security policies and standards to support AI system risk assessments, establish guardrails that enforce risk decisions, facilitate continuous authorizations, and implement new or emerging controls consistent with EO14110, in the National Institute of Standards and Technology (NIST) cybersecurity risk framework, and the NIST AI risk management framework.

AI Talent

The Board's human capital office is engaged with and supports the AI Program to implement best practices in talent selection and staff development for AI. The Board will leverage existing policies, work practices, processes, and resources to provide AI-focused education and development. With collaboration across divisions, and engagement with leadership and staff development experts, the Board will identify and implement AI staffing requirements as well as accompanying educational and development needs.

- **AI talent profiles.** Through leading workforce planning practices, the Board will identify talent profiles needed to develop and update job descriptions and technical competency models to appropriately encompass AI skillsets. Requirements will be documented to analyze job families at all levels and will enable prioritization of AI learning resources, target needs for new talent acquisition, and inform role-based education for existing staff.
- **Internal AI learning and development.** The Board will perform an assessment to determine AI development needs for technical and non-technical personnel. Assessment results will determine appropriate content, modalities, and prioritization. The Board has provided some training on AI and generative AI, which in conjunction with the assessment results, will be a springboard for additional training content design. AI educational content will be developed or sourced

based on assessed needs and Board priorities. In addition, a feedback loop will be utilized to consider emerging needs.

- **Recruitment strategies.** The Board will implement targeted recruitment strategies to attract AI talent, including leveraging and participating in AI-focused job fairs and conferences. Skill-specific job boards and professional association relationships will be utilized to develop talent pipelines for AI and AI-related vacancies, as will engaging with academic institutions with programs developing related student skillsets. The talent acquisition team will leverage the Board's hiring authorities consistent with the Federal Reserve Act to ensure competitiveness and agility throughout the sourcing and hiring process.

AI Sharing and Collaboration

The Board recognizes the importance of collaboration and knowledge sharing in advancing responsible AI innovation. The AI Program efforts in this area include

- **Custom-developed AI code.** Ensuring that custom-developed AI code, including models and model weights, is shared consistent with section 4(d) of M-24-10. The AI Use Case Inventory will be reviewed by the AI Program team to identify use cases with custom-developed AI code, models, and data, and consider sharing with the public.
- **Incentivizing sharing.** Encouraging the sharing of AI code, models, and data with the public and other agencies by providing incentives and support for such initiatives, as appropriate. Identify opportunities for collaboration with other agencies, the governmentwide AI Community of Practice, and other central banks.
- **Coordination efforts.** Ensuring that best practices in the use of AI are disseminated and adopted across the entire Federal Reserve System and that relevant policies and processes (including data access, acquisitions, enterprise architecture, and cybersecurity) are reviewed to ensure alignment in support of AI code and model sharing, disclosure, and peer review.

Harmonization of AI Requirements

To ensure a consistent and unified approach to AI governance, innovation, and risk management, the AI Program has taken steps to harmonize AI requirements.

- **Documentation of best practices.** Documenting and sharing best practices regarding AI governance, innovation, and risk management to ensure that they are applied consistently. Offering employees guides, tools, tips, and lessons learned shared via the AI Program intranet site as well as through the Board AI Community of Practice. Hosting external speakers from academia, other agencies, and the private sector that can inform staff about advances in AI.
- **Interagency coordination.** Engaging in interagency communication efforts to share AI strategies and policy with other federal financial regulatory agencies as well as promoting a community environment with a collaborative approach to AI use.

- **Continuous program maturation.** Updating its AI practices, policy, and guidance, as needed, to reflect emerging trends, risks and threat vectors, technological advancements, and evolving regulatory requirements.

Managing Risks from the Use of Artificial Intelligence

The Board recognizes the value of a comprehensive enterprise risk-management approach to ensure safe and responsible AI innovation.

Determining Which AI Use Is Presumed to Be Safety- or Rights-Impacting

The Board has implemented its enterprise-wide AI policy and corresponding review process to determine which current or planned AI use cases that are determined to be safety- or rights-impacting.

- **Review process.** Each current or planned AI use case undergoes a thorough review and assessment by the CAIO and the AI Program team to determine whether the use case meets the definition of safety- or rights-impacting AI as defined in section 6 of OMB M-24-10.
- **Criteria for assessment.** Our assessment criteria are based on the definitions of safety- and rights-impacting AI and examples of AI presumed to be safety- or rights-impacting in OMB M-24-10 section 6 and Appendix I, respectively. These criteria include whether the AI output would serve as a principal basis for a decision or action and real-world considerations of potential harms to protected or otherwise critical populations, entities, and resources.
- **Supplementary criteria.** The Board may incorporate additional review criteria to assess safety- and rights-impacting AI considerations in response to internal or external developments.

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

- **AI policy and review process.** The Board's AI policy and review process prohibit any use of AI considered to be safety- or rights-impacting without the CAIO's approval, waiver of one or more risk-management practices, or approved OMB extension, to meet risk-management requirements. All safety- or rights-impacting AI use cases undergo a comprehensive risk impact assessment including validation of all risk-management practices defined in OMB M-24-10 section 5(iv).
- **Enforceability and penalties.** Unauthorized or improper use of AI may result in loss of, or limitations on, the use of Board IT resources and in disciplinary or other action, which could include separation from employment.

- **Technical controls.** The Board has technical controls in place to deter, detect, and remediate policy violations. These controls include the ability to terminate instances of non-compliant AI on Board IT resources.
- **Communications and training.** The Board's AI Program team publishes and manages the AI policy through a regularly updated intranet site. The site provides guidance on the AI policy, the process for submitting a use case, and the criteria for determining the permissibility of a use case. The site also offers non-technical and technical AI training materials, a list of best practices for the responsible use of AI, and answers to policy FAQs.

Minimum Risk-Management Practices for Safety- or Rights-Impacting Uses

The Board is implementing a comprehensive environment of controls to encompass the risk-management practices required by OMB M-24-10. The CAIO and AI Program team are responsible for ensuring that these controls are designed and operating effectively to provide sufficient assurance that the Board can mitigate risks from non-compliant AI uses.

- **Impact assessment.** Every AI use case that is presumed to be safety- or rights-impacting undergoes a comprehensive risk impact assessment, which includes a review of controls and processes meeting or exceeding the minimum risk-management practices defined in OMB M-24-10 sections 5(c)(iv) and 5(c)(v). The review process assesses the quality and appropriateness of AI use cases, all data considered for those use cases, purpose of use, and potential harms to health, safety, privacy, security, rights, and opportunities as noted in the Board's criteria for assessment. Considerations for resourcing, security controls, testing, and validation plans are also reviewed.
- **Determination process.** The CAIO, in conjunction with the AI Program team and, as appropriate, senior Board officials, will review whether the AI use case, along with its impact assessment, satisfies the definitions of safety- or rights-impacting in section 6 of OMB M-24-10. The CAIO shall determine whether the AI use case matches the definition of safety- or rights-impacting after considering the conditions and context of the use case and whether the AI is serving as the principal basis for a decision or action.
- **Waiver process.** In limited circumstances, waivers of minimum risk-managements practices may be granted in accordance with OMB M-24-10 section 5(c)(iii). The AI Program will develop criteria to guide consistent decisionmaking for the CAIO to waive risk-management practices, ensuring that waivers are granted only when necessary. Any decisions to grant or revoke a waiver will require documentation of the scope, justification, and supporting evidence. The AI Program team will establish procedures for issuing, denying, and revoking waivers, with oversight by the CAIO and the AI Enablement Working Group.
- **Documentation and validation.** The CAIO is responsible for documenting and validating that current and planned risk-management practices for all safety- and rights-impacting AI use cases are designed and operating effectively. The AI Program team maintains detailed records of all

use cases and extension, waiver, and determination decisions to support consistent reviews, enable effective compliance and reporting, and promote transparency and accountability.

- **Publication and annual certification of waiver and determination actions.** All materials related to a waiver or determination action will be reported to OMB within 30 days. An annual certification process of the ongoing validity of waivers and determinations will be conducted by the CAIO, the AI Program team, and the owners of relevant AI use cases. The AI Program team will develop procedures for certifying all waivers and determinations. A summary of the outcome of the annual certification process, detailing individual waivers and determinations along with justification, will be shared with OMB and the public in accordance with OMB M-24-10 section 5(a)(ii). If there are no active determinations or waivers, that information will be shared with the public and reported to OMB.
- **Implementation and oversight.** The AI Program team has a dedicated workstream with responsibility for the implementation and oversight of risk-management practices. The workstream includes members specializing in relevant mission and compliance functions, including technology, security, privacy, legal, data, and enterprise risk management, and represents a diversity of enterprise perspectives. The group is responsible for promoting consistent and comprehensive AI risk management through the use case review and impact assessment processes. This workstream is also responsible for maintaining a register of enterprise AI risks and associated mitigations to promote active management and accountability across the Board.