



Board of Governors of the Federal Reserve System Compliance Plan for OMB Memorandum M-25-21

September 2025



Overview

In accordance with the Office of Management and Budget's (OMB) Memorandum M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*, the Board of Governors of the Federal Reserve System (Board) is pleased to share its plan for compliance with the requirements of M-25-21.

The Board remains committed to an artificial intelligence (AI) program that seizes the opportunities AI presents with the Board's Chief Artificial Intelligence Officer (CAIO) playing an important role in advancing the Board's AI objectives and promoting AI innovation, adoption, and governance in coordination with appropriate Board officials.

This document describes the Board's efforts to satisfy the requirements of 3(b)(ii) of M-25-21 and section 104(c) and (d) of the AI in Government Act. The Board will report its individual use-case-specific practices in accordance with section 3(b)(v) of M-25-21 separately through the annual AI use case inventory.

Driving AI Innovation

The Board balances innovation with responsibility by implementing appropriate guardrails while removing unnecessary barriers to AI adoption. The Board conducts comprehensive evaluations of the benefits, costs, and risks associated with AI products; these evaluations inform adoption decisions. The Board has also established robust AI governance frameworks, developed experimentation environments, and fostered an organizational culture that supports responsible innovation. In today's rapidly evolving AI landscape, the organization maintains flexibility to take maximum advantage of cutting-edge AI products tailored to valuable use cases. This approach allows the Board to implement governance and technical improvements to enhance and expand staff use of AI across various domains, including writing, coding, and research activities that increase staff efficiency and effectiveness.

Removing Barriers to the Responsible Use of AI

The Board is removing internal barriers to expanding responsible AI use. The following provides a summary of the steps the Board has taken, or plans to take, to mitigate or remove barriers to responsible AI innovation.

- **IT Infrastructure and Data**

- **Secure access to AI computing environments.** The Board grants staff access to internal computing environments equipped to run AI models and internally publishes and renews a list of available, approved AI computing environments. The Board continuously evaluates emerging alternatives for secure AI computing based on new use cases.
- **Secure enterprise AI tools/platforms.** The Board investigates and authorizes 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 AI models and libraries.** The Board permits secure access to open-source AI foundation models and libraries with vetted, acceptable license terms. The Board maintains a list of approved and available models and libraries and evaluates requests for new models on an ongoing basis.
- **Structured pathways from AI pilot to production.** The Board provides technical and governance pathways to transition AI uses from pilot to production, establishing appropriate processes for model validation, code reviews, and security risk assessments.
- **Data accessibility and suitability for AI use.** The Board evaluates whether data are permitted for use with AI and seeks modifications as needed. The Board is exploring ways to facilitate efficient access to unstructured data for use in AI applications.

- **Policy and Operations**

- **Keeping internal AI policy current.** The Board ensures its internal AI policy remains compliant with OMB requirements.
- **Staying current with AI advancements.** The Board monitors emerging AI trends and best practices, regularly evaluates its AI policy for possible updates, and continuously seeks feedback about desired AI capabilities from users.
- **Ensuring value of AI investments.** The Board evaluates and determines investment levels required to successfully manage AI, including technology, governance, operations, incremental cloud compute costs, and skill development, and is developing analytics to measure return on those investments.
- **Building AI knowledge across the organization.** The Board is implementing an expanding range of AI education and training opportunities both to increase organizational adoption and awareness of AI and to upskill its staff.

- **Cybersecurity and Privacy**

- **Ensuring privacy protections in high-impact AI.** The Board ensures that technology and business processes support and sustain privacy protections, civil rights, civil liberties, and to mitigate any unlawful discrimination when using high-impact AI by establishing technical safeguards that restrict access to sensitive information that is not needed for an approved use case and conducting regular audits of AI systems to verify compliance with privacy requirements.
- **Adapting security measures for AI systems.** The Board updates security policies and standards to support AI system risk assessments, establishes guardrails that enforce risk decisions, facilitates continuous authorizations, and implements new or revised control activities consistent with the safety and security priorities articulated in M-25-21. For example, the Board is developing a risk-based governance framework that matches approval authority to risk level; this framework will expand which AI use cases can be reviewed and approved at lower organizational levels as referenced in 3(b)(ii) of M-25-21.
- **Developing oversight procedures.** The Board implements a comprehensive approach combining technological solutions and human expertise to mitigate errors in AI applications. Supervisor and evaluator agents will help detect inaccuracies or fabricated content (“hallucinations”), while human review protocols ensure that AI use cases undergo expert evaluation before use. These “human-in-the-loop” policies establish clear accountability and verification standards, maintaining the Board’s commitment to accuracy and reliability in all AI applications.

AI Sharing and Reuse

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.** The Board will ensure that custom-developed AI code, including models and model weights, is shared consistent with section 2(b)(i) of M-25-21. The Board’s AI Use Case Inventory is reviewed by the AI Program team to identify use cases with custom-developed AI code, models, and data, and they will consider sharing with other agencies and the public, consistent with the Open, Public, Electronic and Necessary (OPEN) Government Data Act and subject to information security restrictions and any applicable licenses or other legal agreements.
- **Coordination efforts.** The Board coordinates with its Federal Reserve System partners along with other federal financial regulatory agencies and members of the federal CAIO Council. The Board shares information with other central banks to promote alignment of policies and processes (including those related to data access, acquisitions, enterprise architecture, and cybersecurity) that support the sharing, disclosure, and peer review of AI models and code.

The resources that could further enable sharing and reuse include

- **Collaboration infrastructure.** A secure knowledge repository with version control capabilities, technical sandboxes for safely testing shared AI models, virtual collaboration workspaces for real-time engagement, and cross-agency portals to facilitate resource sharing with external partners.
- **Governance frameworks.** Resource allocation mechanisms for prioritizing collaborative initiatives with high potential value, regularly enhanced security protocols for managing access to shared resources, and evaluation criteria for assessing the effectiveness of collaboration efforts.

AI Talent

The Board's People, Strategy & Operations office is engaged with and supports the AI Program to implement best practices in talent selection and staff development for AI. The Board leverages existing policies, work practices, processes, and resources to provide AI-focused education and development. The Board offers a diverse learning approach that includes formal training sessions, hands-on workshops, online learning platforms, collaborative brainstorming events, and practical development opportunities, creating a well-rounded curriculum for skill building and knowledge growth. The Board has also implemented several AI training initiatives, including creating and delivering training about permitted uses of generative AI as well as hosting an organization-wide event that allowed teams to ideate and implement AI methods for approved use cases. Furthermore, Boardwide AI training is available to all staff.

- **Internal AI learning and development.** The Board conducted a needs analysis identifying AI-related competency gaps and recommending targeted development aligned with strategic outcomes. The Board has provided training on AI and generative AI, which in conjunction with the needs analysis, serves as a springboard for additional training content design. The Board partnered with Federal Reserve Systems partners to curate and align AI upskilling content, ensuring consistent messaging. The Board developed FAQs and Quick Reference Guides on shared platforms to enhance information accessibility.
- **AI training areas.** The Board identified the AI skills most needed to develop and will update job descriptions and technical competency models to appropriately encompass necessary AI skill-sets. The Board continues to evaluate all existing roles for AI upskilling opportunities. Depending on the role, that may mean, among other things, further developing AI technical skills, expanding the skills needed to evaluate AI outputs, or fostering a greater awareness of AI-related ethical issues and compliance considerations.

- **Recruitment strategies.** When permitted and subject to any internal guidelines, targeted recruitment strategies will be used to attract AI talent, including participation in AI-focused job fairs and conferences. The Board will utilize skill-specific job boards and relationships with professional associations to develop AI-related talent pipelines, while also engaging with academic institutions that develop AI-related skillsets among their students.

Improving AI Governance

AI Governance Board

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. In late 2025, the Office of the Chief Data Officer took over strategic oversight of the AI Program. The AI Program consists of two governing bodies: an AI Program Team and an AI Enablement Working Group that spans the Board's organizational units and key compliance offices. 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. This role includes maintaining the Board's AI use case inventory, evaluating use case permissibility, routing use cases into the appropriate governance path for impact assessment, and making approval decisions. Enterprise AI investment decisions are reviewed by 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 Chief AI Officer, with appropriate agency officials making final decisions. These bodies are critical components of the Board's commitment to promoting AI innovation while ensuring the responsible and ethical use of AI technologies at the Board. Furthermore, to help ensure compliance with internal AI policy, the Board implements technical controls to monitor AI usage, detect policy violations, and terminate noncompliant AI use.

Consultation with External Experts and the Federal Government

The CAIO and AI Program consult with external experts, including academic communities, industry leaders, other federal agencies, and international organizations. For example, Board staff

- collaborate with researchers and experts from universities and research institutions;
- engage with industry experts and vendors to gain insights into AI technologies and services;
- coordinate with other federal agencies, particularly other federal financial regulators, to share knowledge and align best practices for AI governance and operations; and
- share information on AI developments and best practices with other central banks and international organizations in established international forums.

Agency Policies

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 AI Program team manages and publishes the AI policy through a regularly updated intranet site.

The Board considers the following guiding principles for updates to its AI policy:

- Actionable: Enable permissible AI adoption and innovation; don't impede it.
- Comprehensible: Make it easy for users to understand and comply.
- Differentiated: Avoid duplication with other Board policies; focus on what is new and AI-specific.
- Sustainable: Maximize the policy's current utility but minimize the risk of rapid obsolescence.
- Streamlined: Make it simple.
- Compliant: Aligned with applicable and evolving OMB guidance.

The Board's AI policy language is supplemented by additional guidance materials that are maintained by the AI Program team on a dedicated intranet site. These reference materials provide guidance on the AI policy and other compliance requirements related to Generative AI (GenAI) use cases; the GenAI use case submission process; lists of approved models, tools, computing environments, and licensed data inputs for GenAI use; and how to participate in current and planned GenAI projects, pilots, and trials.

AI Use Case Inventory

The creation and maintenance of AI use case inventories are essential to ensuring the Board has a comprehensive understanding of how AI is being adopted across the organization. This inventory process allows the CAIO to manage AI deployments effectively, identify opportunities for enterprise-level innovation, and ensure alignment with the AI policy. For the annual AI Use Case Inventory, the Board conducts the following process for soliciting and collecting AI use cases in Board work and Board-delegated functions executed by the Federal Reserve Banks:

1. The AI Program actively gathers comprehensive AI use case data. All users of AI in connection with Board work document their use cases and submit them to the AI Program.
2. The AI Program reviews each use case to check permissibility under the Board's AI Policy and evaluate exclusion from public reporting as defined by OMB's *Guidance on 2025 Agency Artificial Intelligence Reporting Advancing American AI Act EO 13960*. For covered AI use cases, the AI Program gathers additional detailed information required for submission to the public inventory.

Use cases are stored in a common repository allowing reporting, storage, and ongoing tracking. The AI Program team works with current and future AI users to add new use cases and update existing ones as needed. The Board maintains the AI Use Case Inventory and validates it periodically to ensure accurate and transparent reporting in compliance with OMB reporting requirements.

Fostering Public Trust in Federal Use of AI

Determinations of Presumed High-Impact AI

The Board is implementing a streamlined set of controls to ensure that the Board leverages AI's transformative potential while implementing the necessary controls to comply with the risk-management practices required by OMB M-25-21. Any high-impact AI use requires explicit CAIO approval. Each current or planned AI use case undergoes an initial review by the AI Program team in conjunction with appropriate agency officials using set criteria to screen for high-risk characteristics, including whether the use case appears to meet the definition of high-impact AI per OMB M-25-21. Those use cases flagged as possible high-impact AI uses are referred to the AI Program team for confirmation of the high-impact determination, potential approval, and implementation of the required minimum risk-management practices and/or waivers.

Determination Process

The review process described in “[Determinations of Presumed High-Impact AI](#)” assesses each current and planned use of AI to determine if it matches the definition of high-impact AI. First, the review considers whether the use case is limited to experimental or non-production AI work types (e.g., capabilities testing as well as basic or applied research). If not, then it qualifies as a production or operational use case and *could* therefore serve as the principal basis for a decision or action. Next, the review considers whether the use case matches one or more of the OMB purposes presumed to be high-impact AI as described in M-25-21 Section 6. Finally, the review considers whether the use case meets the definition of high-impact AI.

Impact Assessment

Every AI use case that is determined to be high impact undergoes a risk impact assessment, which includes a review of controls and processes meeting or exceeding the minimum risk-management practices defined in OMB M-25-21. 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 civil rights, civil liberties, safety, or privacy as noted in the Board's criteria for assessment. Considerations for resourcing, security controls, testing, and validation plans are also reviewed.

Waiver Process

In limited circumstances, waivers of minimum risk-management practices may be granted in accordance with OMB M-25-21 section 4(a)(ii). The AI Program team has established procedures for issuing, denying, and revoking waivers to guide consistent decisionmaking for the CAIO. Any decisions to grant or revoke a waiver require documentation of scope, justification, and supporting evidence. Waivers for high-impact use cases are reviewed and recertified annually.

Waiver and Determination Reporting

The AI Program team maintains records of all determinations and waivers, reporting to OMB within 30 days of significant modifications and conducting annual certifications, ensuring continuous improvement while fulfilling transparency requirements. A summary of each determination and waiver, along with its justification, will be shared with OMB and the public in accordance with OMB M-25-21 section 4(a)(iv).

Implementation of Risk-Management Practices and Termination of Noncompliant AI

The CAIO is responsible for documenting and validating that current and planned risk-management practices are designed and operating effectively. The AI Program team maintains detailed records of all use cases and extensions, waivers, and determination decisions. In addition, a risk-management team with an understanding of AI ethics and responsible AI practices helps the organization balance innovation with risk management.

Responsibility of Implementation and Oversight

The AI Program team has a dedicated risk-management team responsible for the implementation and oversight of risk-management practices. This workstream includes members specialized in relevant mission and compliance functions.

Technical Controls

The Board's AI policy and use case review process prohibits any use of AI considered to be high impact without the CAIO's explicit approval. Technical controls are also in place to deter, detect, and remediate policy violations, including anomalous activity monitoring and banners and warnings about impermissible use of AI tools. The Board has technical controls in place to terminate instances of noncompliant AI on Board IT resources. Unauthorized or improper use of AI may result in loss of, or limitations on, the use of Board IT resources and potential disciplinary action.