UNITED STATES OFFICE OF PERSONNEL MANAGEMENT
Washington, DC 20415

Office of Personnel Management

AI Compliance Plan for

OMB Memorandum M-25-21

September 2025

Prepared by: Acting Chief AI Officer Perryn Ashmore
Issued by: Director Scott Kupor

Table of Contents

Overview

The Office of Personnel Management (OPM) is committed to the secure, ethical, and mission-aligned adoption of artificial intelligence (AI) technologies. In accordance with OMB Memorandum M-25-21, this compliance plan outlines how OPM is building the infrastructure, governance, and workforce needed to responsibly deploy AI across the enterprise. This document reflects our current posture and outlines the steps we're taking to ensure AI is used to enhance service delivery, protect individual rights, and earn public trust.

OPM's approach is grounded in transparency, accountability, and innovation. We are building a scalable AI ecosystem that supports experimentation while enforcing guardrails that ensure compliance with federal policy and public expectations. This plan is a living document and will evolve as our capabilities mature and as new guidance emerges.

1. Driving AI Innovation

Removing Barriers to the Responsible Use of AI

OPM has identified several systemic and operational challenges to AI adoption, including inconsistent access to secure development environments, fragmented procurement pathways for AI tools, and limited enterprise-wide visibility into AI activities.

To address these challenges, the Office of the Chief Information Officer (OCIO) has deployed a secure, cloud-native infrastructure built on Microsoft Azure. This environment supports the full AI lifecycle—from sandbox experimentation to production deployment—and includes private endpoints, segmented virtual networks, and role-based access controls.

To democratize access to AI tools, OCIO has aggressively provisioned enterprise-grade generative AI assistants to all knowledge workers. These tools are integrated into the agency's productivity suite and are supported by training, usage guidance, and monitoring to ensure responsible use.

We've also implemented a phased AI development model with clear entry and exit criteria for each stage—sandbox, development, pilot, production, and decommissioning. This model ensures that AI systems are evaluated for risk, compliance, and mission alignment at each step, reducing the likelihood of uncoordinated or non-compliant deployments.

Sharing and Reuse

To promote reuse and reduce duplication, OPM maintains a centralized GitHub Enterprise Cloud repository for AI code, models, and documentation. This repository is managed by OCIO and is accessible to authorized developers across the agency. It includes standardized templates for AI Impact Assessments, model cards, and deployment documentation, which help ensure consistency and traceability.

OCIO also facilitates internal collaboration through regular AI working sessions, technical showcases, and brown-bag briefings. These forums allow program offices to share lessons learned, demonstrate prototypes, and align on best practices. Future plans include the development of a public-facing AI use case catalog to promote transparency and interagency collaboration.

AI Talent

OPM is investing in both the recruitment and development of AI talent. The agency has leveraged government-wide direct hire authority to bring in AI engineers, data scientists, and machine learning specialists. OCIO has also established a dedicated AI engineering team responsible for building and maintaining the agency's AI infrastructure, supporting proof-of-concept development, and advising program offices on technical feasibility.

To build internal capacity, OPM offers role-based training through industry-standard training Initiatives. These programs include learning paths for AI developers, product managers, and non-technical staff. OCIO is also developing an AI Handbook and Playbook to provide practical guidance on topics such as prompt engineering, model evaluation, and ethical considerations. These resources are designed to empower program offices to independently explore and implement AI solutions while maintaining alignment with enterprise standards.

2. Improving AI Governance

AI Governance Board

The AI Governance Board is the central oversight body for all AI activities at OPM. Chaired by the Director, with the CIO/Chief AI Officer serving as Vice-Chair, the board includes the

General Counsel, the Senior Agency Official for Privacy, and other relevant stakeholders. It meets at least quarterly to review AI Impact Assessments, approve progression through development phases, and evaluate waiver requests.

The board's responsibilities include:

- Ensuring that all AI systems are aligned with OPM's mission and strategic priorities
- Reviewing documentation for compliance with privacy, security, and ethical standards
- Coordinating with external experts and interagency partners to stay current on emerging risks and best practices
- Maintaining oversight of the AI use case inventory and risk classification process

The board also plays a key role in fostering a culture of responsible innovation by providing feedback to program offices, identifying opportunities for reuse, and promoting transparency in AI decision-making.

Agency Policies

OPM has updated its internal policies to reflect the requirements of M-25-21. These updates include:

- Revisions to the agency's cybersecurity policy to require AI systems to undergo security impact assessments and obtain Authority to Operate (ATO) prior to deployment
- Updates to the privacy policy to mandate Privacy Threshold Assessments (PTAs) and, where applicable, Privacy Impact Assessments (PIAs) for AI systems that process sensitive or personally identifiable information
- Interim guidance on the use of generative AI, which outlines acceptable use cases, prohibited activities (e.g., uploading sensitive data to public models), and required safeguards

These policies are enforced through the AI Governance Board and are integrated into the agency's broader IT governance and change control processes.

AI Use Case Inventory

OPM maintains a centralized inventory of AI use cases, which is managed by OCIO and

updated as projects progress through the development lifecycle. The inventory includes metadata such as:

- System name and description
- Sponsoring office
- Development phase (sandbox, pilot, production)
- Risk classification (e.g., high-impact, low-impact)
- Compliance documentation (e.g., AI Impact Assessment, PTA, waiver status)

The inventory is reviewed quarterly by the AI Governance Board to ensure completeness and accuracy. It also serves as the basis for OPM's annual reporting to OMB and supports internal planning and resource allocation.

3. Fostering Public Trust in Federal Use of AI

Determinations of Presumed High-Impact AI

OPM uses a structured AI Impact Assessment process to determine whether a system qualifies as high-impact. This determination is based on the following factors:

- An individual or entity's civil rights, civil liberties, or privacy.
- An individual or entity's access to education, housing, insurance, credit, employment, and other programs.
- An individual or entity's access to critical government resources or services.
- Human health and safety.
- Critical infrastructure or public safety.
- Strategic assets or resources, including high-value property and information marked as sensitive or classified by the Federal Government.

The AI Governance Board reviews each assessment and makes a formal determination. Systems classified as high-impact are subject to OPM's risk management practices for high-impact AI use cases as outlined in section 4(b) of OMB M-25-21.

Waiver Criteria and Process

In cases where a program office cannot meet one or more minimum risk management

practices, it may request a waiver from the AI Governance Board. The waiver request must include:

- A justification for the waiver, including technical or operational constraints
- A risk assessment outlining potential impacts and proposed mitigations
- A timeline for achieving full compliance, if applicable

The board evaluates each request on a case-by-case basis and may approve, deny, or conditionally approve the waiver. All waivers are documented in the AI inventory and are subject to periodic review.

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

OPM ensures that all AI systems implement minimum risk management practices through a combination of policy, process, and technical controls. These include:

- Required documentation (AI Impact Assessment, Model Card, PTA)
- Governance checkpoints at each development phase
- Technical safeguards such as access controls, audit logging, and encryption

If a system is found to be non-compliant, the AI Governance Board may suspend its ATO and initiate a decommissioning process. This includes revoking access, archiving data, and updating the AI inventory to reflect the system's status. The board also works with program offices to identify root causes and prevent recurrence.

Conclusion

OPM's compliance plan demonstrates a clear and actionable commitment to the responsible use of AI in alignment with OMB Memorandum M-25-21. Through centralized governance, secure infrastructure, and a growing AI-ready workforce, the agency is building the capacity to scale AI in a way that is ethical, transparent, and aligned with its mission.

As AI technologies evolve, OPM will continue to refine its policies, practices, and oversight mechanisms to ensure that innovation is balanced with accountability and public trust. This plan will be reviewed and updated regularly to reflect new guidance, emerging risks, and lessons learned from implementation.