United States Office of Personnel Management
Washington, DC 20415

OPM.gov

OPM Compliance Plan for OMB Memorandum M-24-10

September 2024

Issued by OPM Acting Director Robert H. Shriver, III

1. Strengthening AI governance

The Office of Personnel Management (OPM) champions human capital for the Federal government and in service to the American people. It has a broad mission to deliver human capital policy, benefits, and services to federal agencies, federal Human Resources (HR) practitioners and leaders, federal employees, and annuitants.

Artificial Intelligence (AI) will enhance the delivery of these functions, enabling greater effectiveness and efficiencies. OPM seeks to utilize AI across the agency in accordance with relevant statute and policy, to include Executive Order 14110 (AI EO) and OMB Memorandum 24-10 (OMB M 24-10).

As required by the AI EO, OPM named an acting Chief AI Officer (CAIO)—its current Chief Information Officer—on December 4, 2023. OPM also updated its Cybersecurity and Privacy policy to state that the use of AI in an OPM system must be coordinated and approved by the OPM AI Governance Board and in accordance with OPM policies and processes governing AI.

In response to OMB M 24-10, OPM established an AI Working Group (AIWG). Overseen by the Acting CAIO, OPM's Senior Advisor for Technology, and OPM's Deputy Chief of Staff, the AIWG has led the agency's efforts to ensure compliance with OMB M 24-10's requirements, including developing an AI governance structure. The AIWG also facilitates the acquisition of the appropriate technical capability, expertise, and talent to promote effective and responsible AI use within the agency and to enhance the agency's enterprise-wide approach.

AI governance bodies

OPM, through the leadership of the AIWG, has established a governance structure to approve and oversee the use of AI at the agency. OPM has established two primary governance bodies: (1) the AI Governance Board and (2) the Chief AI Officer Advisory Group (CAIO Advisory Group). The governance bodies seek to do the following:

- Set and enforce priorities for advancing responsible AI innovation and managing risk for the use of AI to support OPM's strategic plan, IT strategy, data strategy, and other high level guidance documents.
- Govern OPM's use of AI, developing expected benefits while protecting sensitive data, managing risk, and removing barriers to the use of AI.
- Approve the use of data in the AI's design, development, training, testing, and operation in accordance with relevant statute and policy, to include OMB M 24-10.
- Ensure AI is properly coordinated within the agency to advance key agency missions.

AI Governance Board

As contemplated by OMB M 24-10, OPM is taking advantage of an existing structure that the agency had previously developed for data governance—the Data Governance Board—and is using this organizational structure to create the AI Governance Board.

This structure takes advantage of the close relationship between AI and data. Further, OPM already has an IT Strategic Plan and Data Strategy to advance the use of human capital data within the federal government, and continued synergy between OPM's AI efforts and data efforts will be vital to the success of each. Moreover, OPM's full-time CAIO (when hired) will report to OPM's Chief Data Officer (CDO), and both officials reside within OPM's data division, Human Capital Data Management and Modernization.

The AI Governance Board, as required by the AI EO, will be chaired by OPM's Deputy Director (in contrast to the Data Governance Board, which is chaired by the CDO). The AI Governance Board will meet, at minimum, twice a year, as also required by the AI EO. The AI Governance Board membership consists of key agency leaders across a range of functions:

- Deputy Director (Chair)
- Chief Artificial Intelligence Officer (Vice Chair)
- Senior Advisor to the Director for Technology
- Deputy Chief of Staff
- Chief Data Officer
- Chief Management Officer
- Chief Privacy Officer
- General Counsel
- Chief Information Officer
- Director of the Office of Diversity, Equity, Inclusion, and Accessibility (ODEIA)
- Executive Director, HCDMM
- OPM Associate Directors for:
- Workforce Policy and Innovation
    - Retirement Services
    - Healthcare and Insurance
    - Human Resources Solutions
    - Merit System Accountability and Compliance
- Performance Improvement Officer
- Statistical Official
- Evaluation Official

CAIO Advisory Group

The CAIO Advisory Group also performs an important role in OPM's governance infrastructure. The CAIO Advisory Group will conduct the initial review of AI projects at each phase of the product development process. The CAIO Advisory Group, which will be chaired by the CAIO, includes technical experts in data science, artificial intelligence, data engineering, AI/machine learning operations, and product management, as well as experts who can assess risk and make decisions related to the privacy, legal, policy, security, and ethical impacts of AI use cases. It will produce and review documentation, such as the AI Impact Analysis and other required documentation, to properly evaluate each proposed AI project.

AI product development process

The standard product development process includes the following five phases:

1. **Opportunity Evaluation**, which assesses the feasibility, strategic alignment, and potential impact of a proposed AI project;
2. **Research and development**, which is the exploratory research, analysis, and design to develop a testable proof of concept;
3. **Pilot**, which is the continued development and testing of the AI project in a production-like environment to gather feedback on real-world usage;
4. **Production**, which occurs when the AI project is available for use by designated OPM staff, undergoes updates and data refreshes, and is monitored for both cybersecurity risks and resource utilization;
5. **Decommissioning**, which occurs when the AI project is safely retired from use and all associated data and models are shut down.

At inception and throughout the development process, developing safe and effective AI products requires a dedicated development team with a wide-ranging set of skills and expertise, including data scientists, cloud specialists, analysts, designers, data engineers, and domain experts (e.g., HR specialists, personnel research psychologists).

As noted above, AI projects are reviewed and approved at each stage of the development process by the CAIO Advisory Group. When the Opportunity Evaluation and Research and Development phases conclude, the AI Governance Board will evaluate the AI product and must approve any project progressing to the Pilot phase. OPM's CAIO Advisory Group and the AI Governance Board will continue to review the AI product using the established assessment criteria throughout the phases of the development process, which includes a safety-impacting and rights-impacting assessment, including before the AI product transitions into production, and annually during the time it is in production, until it is decommissioned.

The CAIO Advisory Group will have a subgroup consisting of three OPM members: one from the Office of the General Counsel (OGC), one from the Office of the Executive Secretariat and Privacy and Information Management

(OESPIM), and one from the Office of Diversity, Equity, Inclusion, and Accessibility (DEIA). That subgroup will be responsible for drafting a safety-impacting and rights-impacting analysis with a recommendation to the CAIO on whether the proposed AI use case is safety-impacting or rights-impacting.

External experts

Through OPM's AI governance bodies and other fora, OPM will consult with external experts and stakeholders to assist OPM in using AI responsibly and to advance key agency missions. OPM has consulted, and will continue to consult with human capital leaders, as well as labor organizations, to understand how AI will affect federal employees and human capital management within the federal government.

AI use case inventories

As AI projects move from the Research and Development phase to the Pilot phase, they will be added to OPM's AI Use Case Inventory if they meet the criteria for inclusion pursuant to OMB guidance, as determined by the CAIO. Any existing AI use cases that have matured beyond a research and development phase will also be added to OPM's inventory unless it is subject to exclusion.

Any new AI systems being deployed in OPM's environment will trigger an alert to OPM's Security Operations Center (SOC) team. If the newly detected AI system is not registered in the OPM inventory, the SOC will then notify the Chief Information Security Officer, the CIO, and the CAIO of the potential unauthorized risk. This will facilitate timely reviews and approvals, ensuring that all new AI deployments are comprehensively documented and compliant with OPM's AI governance, privacy, and cybersecurity policies.

Further, all AI systems are required to adhere to stringent guidelines to maintain their Authority to Operate (ATO) status, ensuring continuous compliance and oversight. To enhance monitoring, OPM has implemented an IT dashboard that dynamically queries for any AI usage in OPM's cloud environment and displays that usage on a dashboard. This will provide real-time visibility and monitoring of AI activities across all of OPM.

AI use cases not subject to inventory

Throughout the product development and approval process, OPM will assess if an AI product may be excluded from OPM's AI Use Case Inventory. Through a review of the AI Impact Analysis, the CAIO Advisory Group can recommend that the AI Governance Board consider the use case for exclusion from the use case inventory. OPM will follow the criteria for exclusion in OMB guidance and the AI EO for AI reporting. Per EO 14110, an AI product may be excluded from the public-facing inventory if sharing it would be inconsistent with applicable law and governmentwide policy, such as the exemptions from public disclosure provided in 5 U.S.C. § 552. The CAIO Advisory Group's recommendation will document the specific law or policy relied upon to determine that product cannot by shared or released.

OPM will continue to review AI projects and use cases through the product development process. AI products in production will be reviewed annually. AI use cases that were previously excluded from public reporting will be reevaluated prior to the product entering the Production phase and annually thereafter by the CAIO for a determination as to whether it should be added to OPM's published inventory. The criterion for determining exclusion of a use case is based on a determination that sharing would be inconsistent with applicable law and governmentwide policy, such as the exemptions from public disclosure provided in 5 U.S.C. § 552.

2. Advancing responsible AI innovation

Removing barriers to the responsible use of AI

One barrier is OPM's limited access to necessary software tools and open-source libraries. To address this, OPM is making all OCIO approved software tools and open-source libraries available to AI engineers and developers. This includes provisioning licenses for commercial AI tools and setting up a centralized GitHub Enterprise Cloud repository where all approved software tools and libraries can be accessed, promoting consistency and compliance with security protocols.

Security concerns also pose significant challenges. AI systems must be secure to prevent unauthorized data access. To mitigate these risks, OPM has launched the

"Secure Engineering AI Solutions" project, which focuses on developing and implementing secure AI solutions utilizing private networking and zero-trust principles. During the evaluation and research and development phases all AI systems are kept internal to avoid direct exposure to the Internet or external users, creating a secure and controlled deployment environment.

Deployment and monitoring capabilities are crucial for safeguarding AI applications. OPM plans to build comprehensive monitoring for AI at the enterprise level. OPM engineers will develop monitoring capabilities to scan and evaluate all AI solutions from a central console. Additionally, cloud AI threat protection, coupled with OPM cloud AI safety and threat intelligence detection capabilities, will be leveraged to deliver actionable security alerts associated with various threats such as sensitive data leakage, data poisoning, jailbreak, and credential theft.

As to AI use cases requiring integration with existing systems, all such use cases will follow the product development process describe above, such that OPM's governance structure is aware of and signing off on any such cases.

By addressing these barriers through a structured and comprehensive approach, OPM aims to promote the responsible and effective use of AI technologies. This will not only enhance operational efficiency and service delivery but also promote adherence to ethical guidelines and federal regulations, fostering a secure, transparent, and accountable AI ecosystem.

Safeguards and oversight for generative AI

OPM is in the process of developing guidance on the responsible and secure application of generative AI internally and aims to provide a framework that balances innovation with safety and ethical considerations, as well as compliance with relevant statute, regulation, and policy. The guidance will be informed by OPM's previously issued guidance, "Responsible Use of Generative Artificial Intelligence for the Federal Workforce," which OPM issued to the Federal workforce as part of the AI EO.

OPM's AI Governance Board will approve this guidance and will oversee the development, deployment, and use of generative AI within the agency.

Additionally, OPM is implementing detailed security measures aligned with zero trust principles to safeguard AI systems from vulnerabilities and unauthorized access. OPM is developing a comprehensive monitoring system to track the performance and usage of generative AI models, which includes real-time monitoring and periodic audits to identify and promptly address any potential issues. OPM is also establishing safeguards for the use of generative AI applications within the agency. Robust access control mechanisms are in place to control access to AI models and datasets.

AI talent

OPM is the human capital leader of the federal government and understands that the success of its initiatives depends on ensuring that OPM has the right people with the right skills to use AI to achieve key agency missions and goals.

OPM is currently developing its staffing plans for Fiscal Year 2025 (FY 25). Through this process, OPM offices are identifying projected AI roles that they intend to hire in FY 25.

In ensuring the agency hires the right people with the right skills to leverage AI, OPM will build on its own work supporting federal agencies to bring on AI talent. In 2023 and 2024, OPM:

- Authorized a government-wide direct hire authority for Information Technology Specialists; Computer Scientist (Artificial Intelligence); Computer Engineers (Artificial Intelligence); and Management and Program Analysts in the grade levels GS-9 through GS-15. In 2023, OPM also amended government-wide hiring authority for STEM positions in 2023 by adding two occupational series, data science and operations research in the grade levels GS-11 through GS-15 to support federal agency efforts to expand AI capabilities in the federal government.
- Issued Skills-Based Hiring Guidance and Competency Model for Artificial Intelligence, to help agencies identify key skills and competencies that its workforce needs to utilize AI.
- Issued classification policy and talent acquisition guidance to assist agencies in position classification, job evaluation, qualification, and assessments for AI roles

OPM intends to leverage the government-wide direct hire authority for Information Technology Specialists; Computer Scientist (Artificial Intelligence); Computer Engineers (Artificial Intelligence); and Management and Program Analysts in the grade levels GS-9 through GS-15.

OPM will also leverage its AI and Tech Talent playbook, which it will publish by the end of this calendar year. This playbook consolidates federal resources from OPM, the Office of Management and Budget, and others to allow agencies to recruit and effectively utilize AI and Tech talent.

To increase the skills of its workforce, OPM is providing AI training to its workforce. This includes "Gov2Gov" training on AI fundamentals that OPM developed for government-wide use. It will also make available training on AI to its employees, providing them on-demand access to courses, learning journeys, and channels with AI related content for employees at various levels of expertise. OPM will also provide a learning path for in AI for its employees to become AI Certified across a number of AI functions.

In recruiting and developing skills and skillsets for its workforce, OPM will focus on data science and machine learning engineers. As an agency focused on policy delivery and oversight, and benefits delivery, OPM will also focus on AI enabling roles and skills. The agency will target skillsets including data science, data visualization, machine learning engineering, and AI ethics. These positions will be reviewed and tracked by the AI Talent Lead assigned to the OPM Human Resources Office.

Finally, OPM will continue to build out AI-focused teams, embedding knowledge and expertise within OPM organizations, including its Office of the Chief Information Officer, which has teams on AI product development and building enterprise capacity.

AI sharing and collaboration

OPM is committed to fostering an environment of transparency and innovation through the open sharing of AI code, models, and data with the public. This initiative is spearheaded by OPM's Human Capital Data Management and Modernization (HCDMM) directorate, which plays a pivotal role in coordinating and overseeing these efforts.

HCDMM is responsible for developing and enforcing policies that mandate the sharing of AI code, models, and data when legally permissible. These policies are designed to promote the sharing process aligns with the broader goals of transparency and innovation. To support these efforts, OPM will conduct regular training sessions and workshops to educate staff and collaborators on the importance of sharing AI resources. These sessions will cover best practices for data management, code documentation, and compliance with legal and regulatory requirements.

Any custom developed AI code, that conforms to Section 4(d) of M-24-10, will be shared with the public. Additionally, HCDMM actively collaborates with other government agencies, academic institutions, and industry to promote the exchange of AI resources. These relationships help to expand the reach and impact of shared data and models, fostering a broader community of innovation.

The Office of the Chief Information Officer (OCIO), OGC, and OESPIM work closely with HCDMM to align all shared resources with legal, ethical, privacy, and security standards. Together, these offices form a cohesive framework that supports the goal of promoting open sharing while safeguarding sensitive information.

Through these concerted efforts, OPM aims to create a culture of openness and collaboration, ultimately driving advancements in AI technology and its applications for the public good.

Harmonization of artificial intelligence requirements

AI is inherently a crosscutting function involving many different elements of the agency, from data and IT to privacy and legal. Ensuring cross-agency collaboration is vital to the success of AI within OPM.

OPM is developing specific technical guidelines on how to implement an approved use case, to allow OPM to provide cybersecurity and operational consistency across its use cases. OPM will also be developing additional instructional materials for its offices to educate them on the capabilities of AI models and the use cases that other federal agencies have found to be successful. This will help provide a consistent methodology to implement AI use cases, while also educating OPM offices on the possibilities of AI to improve mission delivery.

As OPM continues to refine its AI approval process, OPM will ensure requirements and processes are embedded within agency functions and are well understood within each OPM organization.

3. Managing risks from the use of artificial intelligence

As noted above, a subset of the CAIO Advisory Group—consisting of representatives from OGC, OESPIM, and ODEIA—makes the initial rights-impacting or safety-impacting determination for AI projects during the Opportunity Evaluation phase of the product development process. During this phase, OPM organizations propose AI projects by completing the initial section of the OPM AI Impact Assessment, including questions related to the purpose of the AI, expected benefits and risks, anticipated users, and data to be used. The subgroup of the CAIO Advisory Group will use this information to make an initial rights-impacting or safety-impacting recommendation to the CAIO based on the guidelines provided in OMB M-24-10.

At each of the subsequent stages of the development process, this subgroup of the CAIO Advisory Group reviews the rights-impacting or safety-impacting determination to assess if changes should be recommended to the CAIO based on updates to the AI. Additionally, OPM has supplemented the guidelines of OMB M-24-10 with additional rights-impacting and safety-impacting criteria. These additional criteria are designed to address OPM mission specific uses of AI and clarify for OPM offices that particular uses are right/safety impacting.

OPM program offices can request specific waivers to minimum risk management practices for rights or safety-impacting AI use cases and provide documentation supporting their request in the AI Impact Assessment, which is developed and submitted for approval at each phase of OPM’s AI product development process. The CAIO Advisory Group will review the documentation, and the CAIO will make the final determination for issuing, denying, or revoking a waiver for minimum risk management practices.

OPM has not yet identified any use cases that would require a waiver of minimum practices for rights-impacting or safety-impacting use cases. Prior to waiving any minimum practices for rights impacting or safety impacting AI use cases, OPM will develop criteria to mitigate any potential risks from the waiver of these minimum practices.

Controls and termination of non-compliant AI

As noted above, OPM has implemented a series of controls to align its safety-impacting or rights-impacting AI systems with established risk management practices before deployment. These controls include comprehensive testing, independent audits, and continuous monitoring. These controls are now being folded into our Security Impact Assessment (SIA) and ATO process. AI projects that have not been approved through our governance process and completed our release process will not be deployed.

In cases where an AI system is found to be non-compliant, OPM will terminate its use. This involves immediate cessation of the system's operations, and if it is an ATO based system, suspension of the system's ATO, followed by a thorough review and mitigation plan to address any identified risks or issues.

Risk management practices

OPM moves all AI projects through the AI product development process, outlined above, which incorporates the minimum risk management practices in the day-to-day work of developing, reviewing, approving, and monitoring AI use cases. In each phase of development, OPM program offices are responsible for completing sections of the AI Impact Analysis to provide evidence supporting progressing to the next phase of development. The CAIO, the CAIO Advisory Group, and the AI Governance Board are responsible for reviewing the information provided in the AI Impact Analysis and validating the required minimum risk management practices have been implemented. This includes a routine review of the rights or safety impacting determination for all AI projects and risk mitigation plans.

Additionally, every system employing AI will have its AI functions thoroughly documented as part of its ATO and SIA, which must be completed before any new AI system can be implemented in production.