



U.S. National Science Foundation

AI COMPLIANCE PLAN FOR OMB MEMORANDUM M-25-21 (SEPTEMBER 2025)

Prepared by:

Thu Williams, Acting Chief AI Officer

Issued by:

Simon Malcomber, Chief Science Officer

Micah Cheatham, Chief Management Officer

OVERVIEW

The AI in Government Act of 2020 and U.S. Office of Management and Budget (OMB) Memorandum M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*, direct each agency to submit to OMB and post publicly on the agency's website a plan to achieve consistency with M-25-21, unless the agency makes and posts a written determination that the agency does not use and does not anticipate using AI.

This Plan satisfies the requirements of Section 3(b)(ii) of the Appendix to M-25-21 and Section 104(c) of the AI in Government Act. The agency will report on compliance with the case-specific practices mandated in Section 4 of the M-25-21 Appendix separately, through its annual AI use case inventory.

SECTION 1. DRIVING AI INNOVATION

A. Removing Barriers to the Responsible Use of AI

Describe any barriers to your agency's responsible use of AI, and any steps your agency has taken (or plans to take) to mitigate or remove these identified barriers.

- NSF has identified several barriers to the responsible use of AI, including limited access to AI-ready data, the need to update and streamline AI governance and risk oversight, gaps in workforce readiness, and resource constraints that hinder accelerating responsible AI. To address these challenges, NSF is implementing an internal NSF AI Strategy focused on strengthening AI governance, enabling mission-driven use cases, enabling data and infrastructure for AI, accelerating responsible innovation, and developing an AI-ready workforce. Efforts include updating internal AI policies and use case review procedures, assessing AI maturity, and building centralized cloud-based platforms to provide discoverable, reusable, and scalable AI-ready data and models. NSF is also leveraging interagency partnerships and shared services to reduce duplication and access emerging AI technologies more efficiently. Workforce readiness is being addressed through multi-tiered upskilling, targeted training, and a recently completed AI survey that will guide future investments and capability development. These coordinated steps are designed to remove unnecessary barriers and promote AI innovation in order to drive responsible AI adoption across the agency.

B. Sharing and Reuse

Explain how your agency coordinates internally to promote the sharing and reuse of AI code, models, and data assets. Describe the resources needed to further enable this type of activity.

- NSF plans to promote the sharing and reuse of AI code, models, and data assets through its internal governance structures, available tools, and enterprise coordination mechanisms. The AI Governance Board (AIGB, see Sec. 2.A. below) establishes policies, intake processes, and prioritization frameworks that encourage reuse, reduce duplication, and align AI projects to mission needs. To further enable this activity, NSF plans to explore standing up an internal

repository or AI collaboration platform to store and share reusable code, models, documentation, and best practices. Communities of practice will play a key role in sharing implementation lessons, surfacing reusable assets, and fostering knowledge exchange across staff.

C. AI Talent

Describe any planned or in-progress initiatives from your agency to enhance AI talent. In particular, identify the AI skillsets needed at your agency and where individuals with technical talent could have the most impact.

- NSF has provisionally identified the AI-related technical and non-technical skillsets most critical to accelerating responsible AI innovation, adoption, and deployment across the agency. To address these needs, NSF intends to attract top AI, data science, and cybersecurity professionals through skills-based hiring, pooled certificates, fellowships, and flexible mechanisms such as Intergovernmental Personnel Act (IPA) and Schedule A. Defining AI-focused job roles and skills matrices aligned with NSF use cases will guide recruitment, workforce planning, and development. NSF plans to adopt the annual AI training being developed by OMB and GSA and will supplement this with a multi-tiered internal learning framework. NSF encourages the use of Individual Development Plans (IDPs) to align AI-related learning with mission needs, allowing employees to pursue development goals ranging from foundational knowledge to applied project experience or advanced certifications. To support these efforts, the Chief AI Officer (CAIO) will partner with its internal NSF Academy and Human Resources to track training participation, assess skill gaps, and evaluate metrics to inform future workforce development strategies.

SECTION 2. IMPROVING AI GOVERNANCE

A. AI Governance Board

Describe your agency's AI governance body and the agency offices represented on that body. In particular, describe your agency's plan to achieve the governance board's expected outcomes and describe how, if at all, the governance board plans to consult with external experts and across the Federal Government.

- NSF has established a two-tiered AI Governance ecosystem comprising three governing bodies: an AI Governance Board (AIGB), supported by an AI and Data Committee (AIDC) and an AI Steering Committee (AISC). Both committees report to the Board. The tiered structure for internal AI governance ensures rigorous oversight, transparency and accountability across the entire AI lifecycle. The **AI Governance Board (AIGB)** is responsible for providing strategic direction and oversight of the full suite of agency AI investments for both internal AI use and external NSF research funding and communications. The AIGB approves and oversees implementation of the agency's AI Strategy for integrating AI into agency operations and activities under M-25-21 and serves as the ultimate policy decision-making body for agency AI investments. In alignment with M-25-21, members of the AIGB are:

Core Members:

- Deputy NSF Director (or official performing the duties of the Deputy NSF Director) (Chair) – *position vacant*
- Chief Management Officer (Co-Chair)
- Chief Science Officer (Co-Chair)
- Chief of Staff
- Chief AI Officer (CAIO) (Vice Chair)
- Chief Information Officer
- General Counsel
- Office Head, Award Management
- Asst. Director, Directorate for Computer & Information Science & Engineering
- Asst. Director, Directorate for Technology, Innovation & Partnerships
- Senior Procurement Executive
- Chief Human Capital Officer

Attendance as needed:

- Chief Financial Officer
 - Head, Office of Legislative and Public Affairs
 - Head, Office of Integrative Activities
 - Head, Office of Civil Rights
 - Head, National Center for Science and Engineering Statistics
- The **AI and Data Committee (AIDC)**, chaired by the Chair AI Officer (CAIO), Chief AI Officer (CAIO), will serve as the operational body supporting the AIGB by preparing recommendations, coordinating implementation of AIGB decisions, and ensuring alignment with applicable OMB memoranda. The Committee will provide the AIGB with technical, policy, and operational input in advance of Board deliberations, and will report regularly on progress, risks, and outcomes to inform AIGB oversight. The Committee is responsible for preparing, coordinating, and ensuring timely implementation of the following OMB Memoranda for NSF's internal use of AI: M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*; M-25-22, *Driving Efficient Acquisition of Artificial Intelligence in Government*; and M-25-05, *Phase 2 Implementation of the Foundations for Evidence-Based Policymaking Act of 2018: Open Government Data Access and Management Guidance*. The committee also consults with NSF data officials to align data management, stewardship, and accessibility efforts with the needs of AI systems, ensuring the availability of AI-ready data. The AIDC includes senior agency officials designated under the Evidence Act, as well as key agency advisors or their designees who support the integration of AI technologies into NSF operations in accordance with federal guidance. The members of the AIDC are:
- Chief AI Officer (CAIO) (Chair)
 - Chief Evaluation Officer (EO)
 - Chief Statistical Official (SO)
 - Chief Data Officer (CDO)

- Performance Improvement Officer
 - Office of General Counsel representative(s) (i.e., Assistant General Counsel(s))
 - Data and Analytics Officer representative(s)
 - Office of the Chief Information Officer (OCIO) representative(s)
- The **AI Steering Committee (AISC)** is not directly involved in the development or deployment of AI systems for internal agency use as described in OMB Memorandum M-25-21. Instead, the AISC serves as a cross-agency forum to promote coordination and knowledge-sharing among representatives from NSF Directorates and Offices responsible for awarding and administering grants and cooperative agreements that support AI-related extramural scientific research, research infrastructure, and education, as authorized by the NSF Act and related authorities. The AISC also coordinates NSF’s implementation of executive actions and Congressional direction related to externally focused AI funding initiatives. It keeps members and the AIGB informed of significant AI developments and opportunities—both internal and external to NSF—including those arising from other government agencies, academia, and industry. In addition, the AISC supports responses to inquiries and reporting requirements from Congress, the President, and other stakeholders concerning NSF’s AI research and education funding programs, priorities, and policies.

B. Agency Policies

Describe any planned or current efforts within your agency to update any existing internal AI principles, guidelines, or policy to ensure consistency with M-25-21. Identify whether your agency has developed (or is in the process of developing) internal guidance for the use of generative AI.

- NSF developed initial agency AI policy and procedures in 2023 and has convened ad hoc working groups, led by the Acting CAIO, to update these materials to ensure consistency with M-25-21. The policies and procedures are intended to apply to all internal AI systems in the design, development, acquisition, or deployment stages. NSF will ensure that these policies and procedures appropriately address IT infrastructure, data, cybersecurity, and privacy to remove bottlenecks to AI innovation.
- NSF issued a notice to the research community (Dec. 14, 2023) that prohibits individuals whom NSF has retained as “peer reviewers” from using unapproved generative AI to prepare written reviews of grant proposals for NSF, due to privacy and confidentiality issues arising from such use. The agency intends to clarify this policy and formulate additional policies and guidance on the use of generative AI tools in the conduct of other agency business. All generative AI policies and procedures will include adequate safeguards and oversight mechanisms to allow use without posing undue risks and will prescribe minimum risk management practices for any generative AI uses that qualify as “high impact” under M-25-21.

C. AI Use Case Inventory

Describe your agency's process for soliciting and collecting AI use cases across all subagencies, components, or bureaus for the inventory. In particular, address how your agency plans to ensure the inventory is comprehensive, complete, and encompasses updates to existing use cases.

- NSF intends to issue formal guidance requiring all directorates and offices to submit AI use cases for inclusion in a centralized inventory, overseen by the CAIO and governance bodies. The standardized intake form will capture required metadata aligned with M-25-21, including development stage, intended use, data and model details, and risk mitigation practices.
- Submitters are responsible for maintaining accurate and timely information and must report updates such as deployment status changes, model revisions, or system retirement. The CAIO team will conduct quality assurance, follow up to address gaps, and periodically audit the inventory to ensure completeness and currency.
- This process helps ensure that NSF maintains a comprehensive view of all enterprise AI activities, enabling oversight, reducing duplication, and supporting responsible scaling.

SECTION 3. FOSTERING PUBLIC TRUST IN FEDERAL USE OF AI

A. Determinations of Presumed High-Impact AI

Explain your agency's process to determine which AI use cases are high impact. If the 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. Describe your agency's process for issuing, denying, revoking, certifying, and tracking waivers for one or more of the minimum risk management practices.

- NSF has implemented a standardized AI use case review process, led by the CAIO AI Review Team and AI Risk Management Team, to assess all proposed and existing AI activities. Each use case is submitted through a structured intake form that collects metadata aligned to M-25-21 requirements. During the review process, use cases are evaluated to determine whether they meet the definition of "high-impact AI," based on potential to significantly affect rights, safety, access to services, or other federally defined criteria. Categories presumed to be high-impact are examined closely, and any determinations that a use case does not meet the high-impact threshold are documented and retained by the CAIO for potential audit.
- NSF has not developed supplemental criteria beyond M-25-21 to classify high-impact AI, and no high-impact AI use cases have been identified to date. Use cases are reviewed annually, including reassessment of risk and impact level.
- For waiver management, NSF is developing a standardized waiver template to be used for any exemption from minimum risk management practices. The template will capture system-specific risk justification, supporting evidence, and formal certification. Waivers must be approved by the CAIO, are indexed and stored by use case, and will be tracked for annual validation, reassessment, or revocation if risks increase. All records will be available for OMB review upon request.

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

Identify how your agency plans to document and validate implementation of the minimum risk management practices. Elaborate on the controls your agency has put in place to prevent non-

compliant high impact AI from being deployed to the public. Describe your agency's intended process to terminate, and effectuate that termination of, any non-compliant AI.

- NSF will require all high-impact AI use cases to include a detailed Risk Management Plan at submission, documenting how each of the seven minimum risk management practices defined in M-25-21 will be implemented and maintained throughout the AI system's lifecycle:
 - Pre-deployment testing and evaluation
 - AI impact assessments
 - Ongoing monitoring of performance and adverse impacts
 - Adequate human training and assessment
 - Human oversight, intervention, and accountability
 - Processes for remedies and appeals
 - Incorporation of end-user and public feedback
- The CAIO-led AI Risk Management Team will evaluate each plan for completeness, implementation quality, and alignment with M-25-21 requirements. All documentation will be retained for audit, oversight, and future reporting. Submitters will also be responsible for periodic compliance updates, and plans will be re-evaluated annually.
- To prevent the deployment of non-compliant high-impact AI systems:
 - Only AI systems with approved risk management plans will be granted production access.
 - Strict access controls and deployment gating mechanisms will be enforced.
 - Systems flagged as non-compliant during review or monitoring will be barred from deployment or have their operation suspended until risks are mitigated.
- If a deployed AI system is later found to be non-compliant, NSF will initiate a structured termination process:
 - Immediate suspension of the system's use
 - Notification of the CAIO and governance bodies
 - Documentation of findings and root cause
 - Updates to AI policy or controls, if needed, to prevent recurrence
- This process ensures that only AI systems meet all minimum risk management requirements are used, and that NSF maintains full control over the responsible operation of high-impact AI.