



U.S. National Science Foundation

NATIONAL SCIENCE FOUNDATION (NSF)

AI STRATEGY FOR OMB MEMORANDUM M-25-21 – SEPTEMBER 2025

Prepared by:

Thu Williams, Acting Chief AI Officer

Issued by:

Micah Cheatham, Chief Management Officer

Simon Malcomber, Chief Science Officer



Table of Contents

Executive Summary	3
Introduction	5
NSF AI Vision & Ambition	6
NSF's AI Maturity: A Baseline Assessment	7
AI Strategic Pillars	8
Pillar 1: Governance	9
Oversight & Accountability	9
Policies & Risk Management	10
Track & Plan for AI Investment	10
Pillar 2: Mission-Driven AI Use Cases	12
Alignment with Mission Needs	12
High-Value AI Projects	13
Scaling Successful Pilots into Enterprise Solutions	13
Pillar 3: Data & Infrastructure for AI	14
AI-Enabling Data Infrastructure	14
High-Quality & AI-Ready Data Foundation	14
Secure, Reusable AI Pipelines	15
Pillar 4: Accelerate Responsible Innovation	16
Reusable Code & Models	16
AI Platform with Built-In Guardrails	16
AI Automation (Agentic AI & Model Context Protocol)	17
Pillar 5: AI-Ready Workforce	18
AI Literacy & Awareness	18
Community of Practice (CoP)	18
Retain, Attract, & Upskill Talent	18
Conclusion	20
Appendix A: Mapping of Agency's Responses to OMB AI Strategy Template Questions	21



Executive Summary

Artificial Intelligence (AI) is fundamentally changing how government agencies fulfill their missions. The National Science Foundation (NSF) must embrace this opportunity to remain at the forefront of innovation. Aligned to NSF's mission "to promote the progress of science" and advance national welfare, this internal AI strategy charts a proactive course to enhance NSF's operations, decision-making, and service delivery. It reflects federal best practices for responsible AI use that emphasize innovation, governance, and public trust, while positioning AI as a force multiplier to improve efficiency, enhance decision quality, and maximize mission impact.

NSF aims to accelerate AI innovation, adoption, and AI use to enhance business operations and achieve enterprise-wide improvements in AI maturity. By doing so, NSF will enhance its internal capabilities, make intelligent data-informed decisions, better serve stakeholders, and harness AI to accelerate scientific progress, promote national prosperity, and strengthen U.S. global competitiveness in science and innovation.



Governance

Establish clear oversight, embed risk management, and guide strategic investments to ensure responsible, transparent, and high-impact AI use



Mission-Driven AI Use Cases

Identify, prioritize, and scale high-value and mission-aligned AI initiatives across NSF



AI-Ready Data & Infrastructure

Enable secure and scalable environments that ensure data is accessible, high-quality, and AI-ready across the AI lifecycle




Accelerate Responsible Innovation

Accelerate AI experimentation and deployment while embedding safeguards to protect privacy and promote trust



AI-Ready Workforce

Build AI fluency across the NSF workforce to equip staff with the skills to confidently adopt and manage AI



This AI strategy positions NSF to accomplish its mission efficiently and effectively. By harnessing AI consistent with these five pillars, NSF seeks to reduce administrative burdens, uncover data-driven insights, and deliver faster, higher-quality services to employees, researchers, and stakeholders. Routine tasks that can be automated or augmented by AI will free staff to focus on high-value, mission-centric work. Decision-making can become more data-informed and timely, whether in grant funding analysis, program management, or operational support. The embedded emphasis on governance and responsible AI standards to ensure that these advances occur in a manner that preserves public trust and aligns with our NSF values.

At its core, NSF's enterprise-wide adoption of AI is designed to amplify its capacity to support the science and engineering community, drive greater operational excellence, and exemplify leadership in responsible innovation within the federal government. **Embracing AI in this mission-oriented, principled way will help NSF deliver even greater value to the public and scientific community in the years ahead, ensuring the agency remains a model for modernization and impact in the AI era.**



*Embracing AI in this mission-oriented, principled way
will help NSF deliver even greater value to the public and
scientific community in the years ahead...*



Introduction

NSF's AI strategy is rooted in its unique role as the primary federal funder of both AI research and education. Since the 1960s, NSF has funded research breakthroughs in AI that built the foundation of everyday technologies for the American people. While NSF advances AI research and education, this internal strategy outlines the agency's current state, vision, and approach for harnessing AI to enhance its own operations, decision-making, and mission delivery. This strategy recognizes that while AI offers unprecedented potential to accelerate discovery and address complex challenges, it also raises important concerns related to security, governance, transparency, and trust. The strategy begins by assessing NSF's internal AI maturity and identifying barriers to enterprise-wide AI adoption. It then outlines the agency's approach to enabling responsible and scalable AI through a robust and strategic governance framework, mission-aligned use cases, modern infrastructure, and an AI-ready workforce. Throughout, the strategy emphasizes the importance of operationalizing cutting-edge tools and practices to accelerate innovation, improve decision-making, and enhance NSF's ability to serve the scientific community.



NSF AI Vision & Ambition

NSF envisions an enterprise where AI accelerates mission delivery, enhances operational excellence, and upholds public trust through responsible and secure innovation, deployment, and use. This strategy defines a path to advance NSF's mission by adopting AI tools and methods that are human-centered, responsibly deployed, and operationally transformative. Through enterprise-wide alignment, we will cultivate an AI-ready workforce, and modernize data and infrastructure, specifically to support scalable and secure AI applications, and embed governance to ensure AI is applied transparently, efficiently, and in service of science and society.

To scale AI responsibly and effectively across the Foundation, NSF plans to strengthen and coordinate several foundational enablers. The following strategic enablers serve as the connective tissue between our ambition and our ability to execute:



Align AI efforts with our mission priorities through strong governance oversight, clear scaling criteria, and outcome-based metrics tied to NSF's strategic goals



Ensure data are AI-ready across the lifecycle to accelerate the development and deployment of mission-aligned AI solutions at scale



Equip staff with targeted training, communities of practice, and role-based skills development tied to evolving mission needs



Establish a formal AI use case lifecycle and governance model, backed by agency leadership and embedded policy, to surface and scale high-value projects, eliminate duplication, and ensure transparency and accountability



Enable cross-directorate visibility through a centralized inventory of AI tools, use cases, and investments to support reuse, reduce redundancy, and track performance



Leverage interagency partnerships, shared services, and federal AI tools to accelerate adoption, avoid redundant AI efforts, and maximize return on investment (ROI)



NSF's AI Maturity: A Baseline Assessment

This snapshot provides an assessment of NSF's current maturity in implementing AI across five key pillars: governance and risk management, mission-driven use cases, data and infrastructure, innovation, and workforce readiness. While this is not a comprehensive maturity model assessment, it offers a baseline to guide progress and inform strategic decisions. This overview reflects insights gathered from agency-wide engagements, including AI survey results and AI pilots. The table below outlines the current state in each area, serving as a foundation for the strategic pillars and implementation actions.

AI Maturity Assessment – As of September 2025

Key Enablers	Current State
Governance	- AI governance bodies are formally established - Governance roles, structures, and policy development initiated - Initial AI risk identification practices in place; formal AI risk framework under development
Mission Driven AI Use-Cases	- Independent AI efforts across NSF have led to siloed experiments without clear path to scale - High interest in AI experimentation to solve agency business cases - Initial AI-use case form developed; ROI metrics not yet required
Data & Infrastructure for AI	- Initial data cataloging and metadata management processes launched - Efforts to improve data quality, and discoverability in progress - Fragmented infrastructure not conducive for AI development - Cloud environments for AI tool provisioning are not readily accessible due to resource constraints and provisioning complexity
Accelerate Responsible Innovation	- Lack of a centralized AI development environment (AI sandbox/ platform) to accelerate innovation/scalable deployment - Lack of AI staffing and funding to accelerate AI development and deployment - Limited cross-agency integration of AI experimentation frameworks
AI Ready Workforce	- Limited AI-related workforce training available; AI training not yet mandatory or role-specific - Established Community of Practice to share and learn AI together - Multiple Program Officers on staff with AI expertise can share their knowledge



AI Strategic Pillars

To accelerate AI adoption across NSF successfully, the agency has identified five strategic pillars that form the foundation for responsible, mission-driven innovation. These pillars focus on establishing effective governance oversight, scaling mission-driven AI cases, ensuring access to high-quality data and modern infrastructure, fostering a culture of secure and transparent experimentation, and cultivating a skilled and AI ready workforce. Each pillar is aligned with federal guidance and tailored to NSF's scientific and operational priorities. Together, they offer a unified roadmap to embed AI across the enterprise in ways that are measurable, accountable, and aligned with national priorities and public trust.





Pillar 1: Governance



Goal Statement:
Establish a comprehensive AI governance model that ensures transparent oversight, effective risk management, and strategic investment planning to accelerate responsible AI innovation while maintaining compliance, accountability, and public trust across NSF's AI lifecycle.

Oversight & Accountability

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 will provide strategic oversight over the agency's integration of AI. The AIGB will prioritize and provide recommendations around AI business needs. The AIGB will have primary responsibility for overseeing development and implementation of the agency's AI strategy and risk management processes. The AIGB will review and vet strategic business needs for AI use cases and oversee high-impact use cases.

The AI and Data Committee, chaired by the 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 AI and Data Committee will provide the AIGB with technical, policy, and operational input in advance of the Board deliberations, and will report regularly on progress, risks, and outcomes to inform AIGB oversight. The Committee members will provide input, recommendations, and subject-matter expertise the inform deliberations, and the CAIO holds final decision-making authority for matters within the Committee's scope that do not require AIGB approval. The Committee will keep the AIGB informed of its decisions and deliberations to ensure transparency and alignment with AIGB strategic goals.

The AI Steering Committee 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.

Governing Bodies

**AI Governance Board
(AIGB)**

provides strategic oversight of all NSF AI efforts



**AI and Data Committee
(AIDC)**

serves as the operational body supporting the AIGB



**AI Steering Committee
(AISC)**

serves as a cross-agency forum to promote coordination and knowledge-sharing



The CAIO is responsible for guiding and tracking progress on the AI Governance Framework. This role entails leading and coordination across NSF to establish processes for managing high-impact AI initiatives. Responsibilities include advancing responsible AI innovation and adoption across NSF, guiding the transformation of the workforce into an AI-ready posture, and overseeing the internal inventory and sharing of custom-developed AI code and data to promote reuse and collaboration. The role also supports efforts to monitor AI-related expenditures and serves as a strategic liaison to internal stakeholders, interagency partners, and the broader AI community, ensuring agency activities align with applicable laws and government-wide guidance. To institutionalize this governance model, NSF is developing a comprehensive AI Governance Framework. This framework will include policy-driven oversight, investment tracking, and mechanisms to align AI activities with NSF's strategic priorities and federal mandates. NSF plans to implement the framework in a phased approach, beginning with pilot programs to refine policies and procedures before scaling to mission critical applications. By institutionalizing these structures and responsibilities, NSF aims to foster governance that accelerates AI innovation while ensuring transparency, accountability, and public trust in the responsible development, deployment, and oversight of AI technologies.

Policies & Risk Management

NSF plans to adopt a systematic approach to identify, monitor, evaluate and mitigate risks across the AI lifecycle, ensuring compliance with federal regulations while upholding the principles of transparency, accountability, privacy, and security in the responsible AI and risk management framework. To support secure innovation and oversight, NSF plans to expand its AI infrastructure with tools such as secure testing environments (e.g., AI sandboxes), validation systems, and post-deployment monitoring. NSF plans to establish personnel policies that define acceptable use of AI in work products, enabling the workforce to responsibly explore AI tools and innovate within well-defined boundaries. These capabilities will help assess model performance, support regulatory compliance, and enable effective incident response.

Effective risk management for high-impact and generative AI use cases requires close collaboration across a cross-functional ecosystem. This includes legal, compliance, data governance, evaluation, privacy, IT security, business units, and mission stakeholders, all to ensure responsible, trustworthy, and mission-aligned AI deployment. Policies will aim to implement traceable decisions and risk-aware system design. Compliance reviews will be supplemented by performance monitoring, privacy assessments, and audits. Continuous improvement will be driven by performance indicators (e.g., process efficiency, risk mitigation outcomes, user experience), incident tracking, root-cause analysis, and implementation of any lessons learned. Feedback mechanisms, including stakeholder surveys, will ensure the governance framework and risk management remains adaptive and aligned with public and organizational expectations.





Track & Plan for AI Investment

NSF plans to implement a structured process to plan, track, and evaluate all agency AI investments. A centralized investment tracker will monitor AI-related funding, tools, infrastructure, and procurements. This system will ensure alignment with NSF's strategic priorities, reduce duplication, and inform data-driven resource allocation. Under the oversight of NSF's governance bodies, investment performance will be assessed using defined metrics such as usability, productivity, reliability, security, and cost to support both transparency and continuous improvement.

Procurement planning will take a proactive and collaborative approach by engaging key stakeholders across the agency. NSF intends to utilize performance-based acquisition strategies for AI acquisitions, including incorporating clear performance thresholds, quality assurance plans, and escalation paths to ensure contract accountability throughout the system's lifecycle. NSF intends to adopt procedures for procuring large language models that undergo content neutrality review. This will allow outputs that comply with unbiased AI principles and are consistent with the agency's scientific integrity and communication standards.

NSF intends to update its acquisition and non-disclosure clauses and to safeguard federal data and explicitly address AI-related risks. Acquisition documents will include appropriate deliverables such as data provenance, model documentation, and system transition plans to preserve interoperability and continuity. NSF's AI investment strategy and acquisition planning will also support NSF's ability to maintain appropriate control of AI training data and derivatives and promote open competition across the federal AI marketplace.

Future investment planning will focus on scaling reusable capabilities, closing critical infrastructure gaps, and enabling secure AI innovation. NSF will prioritize U.S.-based AI solutions, where feasible, in support of national innovation and security goals. Investments will be guided by clear criteria to ensure each AI systems and procurement practices deliver mission-relevant outcomes while maintaining transparency, accountability, and public trust.

NSF will actively engage in the Federal Chief AI Officer Council and collaborate with other federal agencies to align with OMB guidance and advance the priorities established under the Trump Administration's American AI Initiative, working together to develop government-wide scientific AI standards, eliminate duplicative efforts, and leverage cross-agency initiatives to accelerate responsible AI innovation.

Pillar 2: Mission-Driven AI Use Cases



Goal Statement:

Identify, prioritize, and scale high-impact AI use cases that both align with NSF's mission and deliver measurable value. By institutionalizing a clear AI use case lifecycle and governance process, NSF will ensure responsible deployment and maximize return on investment across the enterprise.

Alignment with Mission Needs

NSF intends to establish an enforceable AI use case inventory to include a policy-mandated, centralized AI use case inventory and intake portal. This will create an authoritative list of AI needs and projects, enabling NSF to centrally acquire and provide AI resources that are timely and relevant. This structured management of AI use cases will allow the agency to shift from fragmented experimentation to enterprise-coordinated, mission-aligned AI investments that improve efficiency and reduce duplicative efforts. This approach will also reduce redundant efforts by mandating inventory disclosure and enabling cross-directorate visibility, while enhancing AI return on investment (ROI) through structured, data-driven evaluation and prioritization.

Below are some examples that illustrate the impact of this structured approach and how a centralized AI use case inventory can drive operational efficiency, mission alignment, and strategic resource allocation across the agency.



Resubmit Checker

This application takes a list of proposal IDs and returns a list of similar proposals submitted to NSF in the past, scored by their similarity. This information is valuable for identifying resubmitted proposals for multiple purposes, including assessing potential Return Without Review (RWR) decisions due to insufficient revision, maintaining continuity in the merit review process, and analyzing broader trends in proposal resubmissions.



Now Assist

Through the implemented Now Assist platform, NSF is exploring whether the Generative Artificial Intelligence (GenAI) features of the ServiceNow Platform can improve customer experience by automatically generating content (such as responses, work notes, and knowledgebase articles). These features have the potential to increase productivity and speed of delivery through intelligent recommendations, and to resolve issues swiftly by utilizing chatbots.



AI-Ready RPPR Data Solution

Research Performance Progress Reports (RPPRs) are collections of documents and reports submitted annually by researchers funded by NSF and contain data crucial for tracking the outcomes and defining the impact of NSF-funded research. This pilot solution will use AI techniques such as Large Language Models, entity recognition, key phrase extraction, similarity search, and relationship detection to extract and index non-machine-readable RPPR data to make it more accessible for making informed decisions and assessing NSF's societal impacts.



To foster a culture of continuous improvement and adaptive governance, NSF intends to institutionalize feedback into the governance model to improve decision-making, increase transparency, and promote innovation. We intend to update the AI use case entries annually, at a minimum, to monitor development, assess business impact and investment, and ensure business value. Reviews will be conducted quarterly, evaluating system performance, pilot outcomes, and enterprise adoption trends. The AI use case inventory and dashboard will be used to identify AI gaps, redundancies, and success patterns. Additionally, internal staff will be encouraged to propose AI innovations with traceable performance goals to ensure field input informs strategy.

High-Value AI Projects

All uses of AI at the agency require approval through a standardized AI use case review process. Assessments of each use case will be made by a designated team using an evaluation framework to ensure the use case is consistent with the evaluation criteria such as ROI metrics (e.g., cost/time savings). To support high-value determinations, all pilot programs of limited scale and duration will be required to define key performance indicators and expected benefits upfront. These high-value determinations will also be directed to NSF AI Governance Board for prioritization and resource allocation to promote investment in strategic AI efforts, maximize reuse, and minimize similar AI efforts across directorates.

Scaling Successful Pilots into Enterprise Solutions

To scale AI initiatives responsibly and effectively, NSF will adopt an AI Scaling Pyramid framework that governs the lifecycle of AI projects from pilot to enterprise implementation. This framework will provide a structured path for evaluating and advancing AI solutions based on defined readiness criteria and measurable mission value. Only projects that demonstrate clear alignment with NSF's strategic goals and deliver tangible impact will progress to broader deployment.

To support reliable scaling, NSF will standardize AI deployments using reusable components, automated machine learning operations (MLOps), and processes that ensure repeatable and trustworthy AI outputs. These deployments will be continuously monitored to ensure performance, integrity, and compliance with agency policies.

A formal governance process will oversee the transition from experimentation to operationalization, ensuring that scaling decisions are transparent, risk-informed, and outcomes-driven. NSF will also foster cross-functional collaboration across directorates by leveraging its AI Community of Practice, encouraging shared learning, breaking down silos, and accelerating enterprise-wide adoption of high-value AI solutions.

Pillar 3: Data & Infrastructure for AI



Goal Statement:

Leverage modern data and AI platforms to accelerate the development, deployment, and scaling of secure, mission-aligned AI solutions. By integrating AI-ready data pipelines, reusable infrastructure components, and embedded development, security, and operations (DevSecOps) practices, NSF will enable responsible innovation across the agency.

AI-Enabling Data Infrastructure

To accelerate AI innovation and ensure responsible development at scale, **NSF plans to modernize its data infrastructure by investing in advanced data platforms with scalable processing capabilities.** These platforms can support large-scale model training and reduce computational costs which will form the backbone of a resilient, AI-ready environment. To improve discoverability, transparency, and reproducibility, this modern platform will allow NSF staff access to the same version-controlled information and will mitigate duplication of effort and infrastructure needs. NSF plans to migrate siloed internal datasets into a centralized platform to unlock data access, support integrated analytics, and enable consistent AI development workflows across the agency.

Reducing data processing time is a key priority to accelerate the delivery of AI solutions and improve efficiency by reducing technology costs. By streamlining the data lifecycle and enhancing data-as-a-service capabilities, NSF can build a more efficient and agile data, analytics, and AI ecosystem. This ecosystem will be supported by the implementation of a data catalog tool, which will manage enterprise-wide traceability of both data and model assets. The catalog will enable staff across NSF to discover, share, and reuse critical data resources in a secure and governed manner. The data catalog tool will equip NSF staff in locating relevant data sources, understanding their nuances and context, as well as supporting security by ensuring appropriate levels of access.

NSF will link source data to model outputs through lineage dashboards and audit trails, to the extent permitted by data confidentiality and privacy laws. These capabilities will help track the full lifecycle of AI models and promote transparency in their development and deployment. Together, these infrastructure investments will support a foundation of AI-ready, interoperable, and traceable data that drives scalable, mission-aligned innovation across NSF.

NSF plans to modernize its data infrastructure by investing in advanced data platforms with scalable processing capabilities.

High-Quality & AI-Ready Data Foundation

To fully support the development and scaling of responsible AI, NSF will establish a dual foundation of high-quality and AI-ready data. High-quality data will be ensured through a robust data quality framework, supported by automated pipelines for cleaning, tagging, validating, and monitoring datasets. These pipelines will enforce minimum data standards for completeness, consistency, metadata documentation, and de-identification that will improve model performance metrics while reducing risks of error propagation and bias in downstream models and help ensure appropriate privacy and confidentiality protections.

Simultaneously, NSF will define and operationalize what constitutes AI-ready data by moving beyond conventional quality metrics. AI-ready data must be machine-readable (e.g., structured formats like JSON or Parquet), machine-understandable (e.g., labeled data with defined schema or ontology), and explicitly structured for AI consumption and reuse. It must also be relevant to the AI use case at hand and compliant with FAIR principles (Findable, Accessible, Interoperable, Reusable), which serve as the baseline for scalable, interoperable AI adoption.

To ensure trust and reproducibility, NSF will embed bias audits, traceability, and quality gates within all pipelines. Additionally, feedback loops will be implemented to support ongoing data improvement informed by AI performance and user input. Collectively, these capabilities will enable NSF to scale AI solutions confidently knowing that the underlying data is not just clean and compliant, but also fit-for-purpose, efficient to process, and optimized for responsible AI innovation.

Secure, Reusable AI Pipelines

NSF plans to enable scalable, secure, and efficient AI development. NSF intends to implement automated end-to-end AI-ready pipelines that cover the full lifecycle — from data preparation to model training and deployment. These pipelines will be designed for repeatability, interoperability, and speed, enabling rapid experimentation and reliable scaling across use cases.

DevSecOps practices will be embedded throughout to ensure continuous integration and delivery with built-in security, compliance, and traceability. Standardized, reusable pipeline templates will be developed to streamline onboarding, eliminate duplication of effort, and reduce time-to-deploy. These pipelines will incorporate version control, lineage tracking, and quality gates to support trustworthy, auditable AI outcomes.

This modular and secure pipeline infrastructure will empower cross-functional teams to reuse and adapt components across directorates, accelerating NSF's ability to deliver mission-aligned AI solutions responsibly and efficiently.

Pillar 4: Accelerate Responsible Innovation



Goal Statement:

Accelerate AI innovation by integrating emerging technologies such as agentic AI, Model Context Protocols and a secure AI platform to enable rapid development, automate complex tasks, and achieve greater impact with fewer resources.

Reusable Code & Models

To accelerate the adoption of secure and mission-aligned AI solutions, NSF will plan to prioritize reusable code, models, and tools designed for enterprise-wide reuse from inception. A conceptual “AI Commons” could serve as a centralized, curated workspace for hosting models, code, outputs, labeled data, and documentation, promoting collaboration and efficiency. Through model standardization, lineage tracking, and interagency coordination – including participation in the Federal CAIO Council – NSF could streamline access to pre-vetted, authorized AI assets, reducing redundancy and lowering costs. For example, if another agency has an approved HR analytics tool, NSF could adopt it under a shared services agreement, saving development time while ensuring compliance and security. The agency intends to analyze the results of its internal AI survey to identify where there may be gaps and needs that might be fulfilled using AI. NSF intends to evaluate the performance of such solutions, while continuing to monitor improvements in AI technology for new developments and next-generation solutions.

AI Platform with Built-In Guardrails

To accelerate AI innovation, NSF plans to deploy a unified AI platform that integrates MLOps capabilities and a secure AI Sandbox to support responsible, scalable, and cost-effective AI development. This platform will serve as the central environment for managing the entire AI lifecycle from data ingestion and model training to deployment, monitoring, and retraining while embedding critical safeguards such as built-in governance controls, human-in-the-loop workflows, reproducibility, and audit traceability. The MLOps layer will streamline development through standardized, reusable pipelines, reducing integration overhead and increasing development velocity across teams.

The AI Sandbox, as part of this platform, will offer a secure, access-controlled space for experimentation with approved datasets, enabling innovation while mitigating risk. Together, this platform ecosystem will enhance compliance with federal requirements, improve model reliability and explainability, and foster collaboration across directorates through shared tooling and templates. By reducing duplication, improving lifecycle transparency, and accelerating delivery of mission-driven AI solutions, NSF will strengthen operational efficiency, ensure trust, and maximize the impact of its AI investments by doing more with less and ensuring responsible, scalable innovation at enterprise scale.

AI Automation (Agentic AI & Model Context Protocol)

To further accelerate operational efficiency and reduce manual burden, **NSF plans to adopt agentic AI systems and modern integration protocols that enable autonomous execution of routine and time-consuming tasks with built-in guardrails and human-in-the-loop oversight.** These systems can be built using leading open-source frameworks and models, which support modular, goal-driven agents capable of performing complex workflows across data and system boundaries.

To streamline how AI models access data, systems, and computing environments, NSF intends to implement the Model Context Protocol (MCP), an emerging open standard designed to act as a universal connector between AI models and their operating environment. Akin to a “USB-C for AI applications,” MCP will reduce integration complexity, improve security, and enable greater reuse of AI models across directorates. It will ensure that models are context-aware, permissions-aligned, and operate within predefined governance constraints.

These technologies will support the development of agentic AI systems that autonomously build and execute extract, transform, load (ETL) pipelines, perform data quality checks and bias evaluations, assess model performance and risk, recommend proposal reviewers based on content and prior history, check for conflicts of interest, and flag incomplete or duplicative submissions.

.....

NSF plans to adopt agentic AI systems and modern integration protocols that enable autonomous execution of routine and time-consuming tasks with built-in guardrails and human-in-the-loop oversight.

Pillar 5: AI-Ready Workforce



Goal Statement:

Equip NSF's workforce with the knowledge, skills, and tools to responsibly adopt, manage, and innovate with AI by fostering AI literacy, building technical capacity, and embedding AI competencies across roles and leadership levels.

AI Literacy & Awareness

NSF prioritizes fostering a culture where all employees, regardless of their roles or technical backgrounds, can meaningfully contribute to the agency's advancement in artificial intelligence. To achieve this, NSF has devised a comprehensive tiered learning approach. Resources include curated learning materials available through LearnNSF, NSF's internal platform for learning, free external learning courses, and in-house events. Additionally, NSF plans to integrate AI fundamentals into the onboarding process and annual training for all staff. To simplify complex concepts, "AI 101" sessions, use-case inspired training, and brown bag discussions are regularly held. Furthermore, AI literacy is incorporated into leadership development programs to ensure senior leaders possess a deep understanding of both the potential benefits and risks associated with AI, empowering them to lead with confidence. Staff are also encouraged to pursue verifiable, no-cost AI certifications through trusted providers, equipping them with practical, industry-relevant skills. These efforts ensure all employees can participate confidently in NSF's AI journey, make informed decisions, and model responsible AI use.

Community of Practice (CoP)

NSF will expand its internal communities of practice (CoP) to foster collaboration and knowledge sharing among its staff. These CoPs are designed to drive peer learning and promote the reuse of valuable insights and resources. To facilitate this, NSF will implement several ongoing initiatives. Regular CoP-led events showcase successful AI pilots and deployments within the organization, while an internal AI knowledge Hub serves as a central repository for sandbox policies, evaluation frameworks, regulatory updates, and other reusable resources. Additionally, NSF will leverage its CoPs to embed AI awareness into broader data and AI literacy training, thereby enhancing collective understanding and promoting transparency as AI capabilities expand.



Retain, Attract, & Upskill Talent

NSF is strengthening its talent pipeline to ensure it can recruit strategically, retain high-performing individuals, and build long-term workforce readiness. By embedding AI literacy at every level, aligning talent with mission needs, and fostering a culture of ethical innovation, NSF is building the capacity to lead in AI responsibly and effectively. NSF has preliminarily 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. To incentivize professional growth and signal long-term commitment, NSF plans to embed AI competencies into career ladders and performance plans. The agency also intends to launch dedicated development tracks for key AI roles such as data scientists, AI program leads, and product managers. These planned actions will help to ensure NSF attracts, develops, and retains the talent required to responsibly scale AI across the agency.

NSF will continue to align AI talent development with directorate-specific needs and long-term mission goals. To ensure that NSF's workforce is AI-literate and AI-ready, technical talent will be prioritized for deployment in high-value mission areas such as merit review, grant fraud prevention, cyber- and research security, enterprise service automation, and generative-AI knowledge delivery. NSF plans to build cross-functional AI expertise and expose staff to real-world AI applications by providing rotational experiences and mission-aligned opportunities. The agency also plans to leverage directorate-led AI pilots to upskill staff by doubling as learning environments and incubators for internal champions.

An AI-ready workforce requires more than skills; it requires a culture that promotes ethical use, transparency, collaboration, and responsible experimentation. A culture of continuous learning is the foundation for lasting AI success, and NSF is planning to support this shift by delivering oversight-focused training using checklists, briefings, and scenario-based learning. Developing AI ethics training in coordination with the AI Governance Board, Office of General Counsel, the research community, and the National Center for Science and Engineering Statistics will ensure that the content reflects NSF-specific legal and ethical considerations. NSF plans to create learning environments that encourage cross-role dialogue and experimentation with emphasis on highlighting and recognizing internal champions to reinforce innovation and shared ownership.



Conclusion

NSF's AI strategy outlines a clear path to harness AI for scientific and operational advancement while fostering transparency and trust. By promoting strong governance, workforce readiness, secure infrastructure, and AI-ready data, NSF ensures AI adoption aligns with national priorities and supports responsible innovation. This living strategy will evolve with emerging technologies, risks, and opportunities—reinforcing NSF's commitment to equity, accountability, and mission impact. Through intentional scaling of high-value use cases and cross-agency collaboration, NSF will lead by example in deploying AI that improves decision-making, enhances operational efficiency, and delivers better outcomes for the American public.



Appendix A: Mapping of Agency's Responses to OMB AI Strategy Template Questions

OMB Questions	Document Section Containing Response
Provide examples of significant agency AI use cases currently in use or planned to be in use.	Pillar 2: Mission-Driven AI Use Cases
Provide an assessment of the agency's current state of AI maturity and a plan to achieve the agency's AI maturity goals in the following key areas.	NSF's AI Maturity: A Baseline Assessment
Describe your agency's plan to develop AI-enabling infrastructure across the AI lifecycle including development, testing, deployment, continuous monitoring.	Pillar 3: Data & Infrastructure for AI
Describe your agency's plan to ensure access to quality data for AI and data traceability.	Pillar 3: Data & Infrastructure for AI
Describe your agency's plan to recruit, hire, train, retain, and empower an AI-ready workforce and achieve AI literacy for non-practitioners involved in AI.	Pillar 5: AI-Ready Workforce
Describe your agency's efforts to provide AI tools and capacity to support the agency's AI research and development (R&D) efforts.	Pillar 4: Accelerate Responsible Innovation
Describe your agency's plan to develop enterprise capacity for AI innovation.	Pillar 1: Governance
Describe your agency's plan to develop the necessary operations, governance, and infrastructure to manage risks from the use of AI, including risks related to information security and privacy.	Pillar 1: Governance
Describe your agency's plan to identify, track, and facilitate future AI investment or procurement.	Pillar 1: Governance



U.S. National Science Foundation