

OTHER LINKS



REPORT A SAFETY CONCERN
[[/ABOUT-NRC/REGULATORY/ALLEGATIONS/SAFETY CONCERN.HTML](#)]

SEARCH

[[/SITE-HELP/SEARCH/RESULTS?SITE=ALLSITE](#)]

[<https://twitter.com/nrcgov>]



[<https://www.instagram.com/nrcgov>] [<https://www.facebook.com/nrcgov/>]



[<https://www.youtube.com/user/NRCgov/>]



[<https://www.linkedin.com/company/u-s--nuclear-regulatory-commission/>]



[<https://www.flickr.com/photos/nrcgov/sets>]



[https://service.govdelivery.com/service/multi_subscribe.html?code=USNRC]



[</public-involve/rss-feeds.html>]

NRC AI Compliance Plan



On this page:

- [Purpose \[#purpose\]](#)
- [Authority \[#authority\]](#)
- [Scope \[#scope\]](#)

Purpose

The AI in Government Act of 2020 and OMB Memorandum M-24-10, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence, direct each agency to submit to OMB and post publicly on its website either a plan to achieve consistency with M-24-10 or a written determination that the agency does not use and does not anticipate using covered AI.

This document outlines the minimum information required for the US Nuclear Regulatory Commission's compliance plans that will satisfy the requirements of Section 3(a)(iii) of M-24-10 and Section 104(c) of the AI in Government Act. NRC will report compliance with the individual use-case-specific practices mandated in Section 5(c)(iv) and (v) of M-24-10 separately through the annual AI use case inventory.

Authority

The establishment of AI policies within NRC is primarily guided by mandates from the Office of Management and Budget (OMB), Presidential Directives, and other federal regulations. OMB mandates, such as the Federal Data Strategy, the Cloud Smart Strategy, and M-10-24, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence, provide a framework for leveraging data as a strategic asset and adopting modern technology practices, including AI. Presidential directives and national strategies, such as the Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (Oct 20, 2023), further outline the government's commitment to advancing AI technology for the public good while ensuring ethical, secure, and transparent use. These authorities collectively empower federal agencies to develop and implement AI policies that align with national priorities, promote innovation, and uphold the principles of accountability and fairness in the use of AI technologies.

Scope

This AI Compliance plan applies to NRC, including their employees and all third parties (such as consultants, vendors, and contractors) that use or access any

Information Technology (IT) resources under the administrative responsibility of the NRC or its IT services. This encompasses systems managed or hosted by third parties on behalf of the agency. All organizational units within the agency must abide by the compliance policies outlined in this document.

This policy covers all technology systems that deploy AI technology, hereinafter called "AI systems." AI is a machine-based system that can make predictions, recommendations, or decisions influencing real or virtual environments for a given set of human-defined objectives. AI systems use machine- and human-based inputs to perceive environments, abstract perceptions into models through automated analysis, and use model inference to formulate options for information or action. The definition includes systems using machine learning, large language models, natural language processing, computer vision technologies, and generative AI. Still, it excludes basic calculations, basic automation, or pre-recorded "if this, then that" response systems.

This policy applies to all new and existing AI systems developed, used, or procured by NRC, which could directly impact the mission or security of the NRC. It does not govern regulatory or other actions regarding non-agency uses of AI.

- [Strengthening AI Governance \[#strengthening\]](#)
 - [General \[#general\]](#)
 - [AI Governance Bodies \[#governance-bodies\]](#)
 - [AI Use Case Inventories \[#use-case-inventories\]](#)
 - [Reporting on AI Use Cases Not Subject to Inventory \[#reporting\]](#)
- [Advancing Responsible AI Innovation \[#advancing\]](#)
 - [AI Strategy \[#strategy\]](#)
 - [Removing Barriers to the Responsible Use of AI \[#removing-barriers\]](#)
 - [AI Talent \[#talent\]](#)
 - [AI Sharing and Collaboration \[#sharing\]](#)
 - [Harmonization of Artificial Intelligence Requirements \[#harmonization\]](#)
- [Managing Risks from the use of Artificial Intelligence \[#managing\]](#)
 - [Determining Which Artificial Intelligence Is Presumed to Be Safety-Impacting or Rights-Impacting \[#determining\]](#)
 - [Implementation of Risk Management Practices and Termination of Non-Compliant AI \[#implementation\]](#)
 - [Minimum Risk Management Practices \[#risk-management\]](#)

|  [RETURN TO LIST \[#scope\]](#)

Strengthening AI Governance

General

Planned or current efforts to update any existing internal AI principles, guidelines, or policy to ensure consistency with M-24-10:

- The NRC did not develop any internal AI principles, guidelines, nor policies prior to the issuance of M-24-10 that would require update.

|  RETURN TO LIST [\[#scope\]](#)

AI Governance Bodies

The offices that are represented on the NRC's AI governance body:

- Office of the Chief Information Officer
- Office of Small Business and Civil Rights
- Office of Nuclear Regulatory Research
- Office of the Chief Human Capital Officer
- Office of Administration
- Office of the General Counsel
- Office of the Executive Director of Operations
- Office of Nuclear Reactor Regulation
- Office of Nuclear Materials and Safety and Safeguards
- Office of Nuclear Security and Incident Response
- Regional Offices
- Office of the Inspector General

The expected outcomes for the AI governance body and the NRC's plan to achieve them:

- The NRC AI Governance Board (AIGB) was established to remove barriers to the use of AI and to manage its associated risks by convening senior agency officials to discuss the governance, risks, and benefits of leveraging AI solutions. The primary outcomes of the board include:
 - Senior Leadership Awareness – Board meetings will consist of informational sessions relevant to AI solutions in consideration for agency use. This will ensure that senior leadership have a general understanding of the concepts and technologies to be applied to specific agency functions.
 - Risk-Informed Decision-Making – Board meetings may result in decisions regarding governance of the use of new capabilities and/or decisions regarding capabilities that should be disallowed within the NRC environment. These decisions will leverage risk inputs from the various senior managers with input from functional subject matter experts (SMEs) and/or external experts as appropriate.

Description of NRC plans to consult with external experts as appropriate and consistent with applicable law:

- The Chief AI Officer (CAIO) may choose to engage external experts via local acquisitions or interagency agreements on a variety of AI topics. As appropriate, the CAIO will provide a forum for the AIGB to review expert findings, observations, recommendations, and/or general field expertise.
- The CAIO will continue to engage in relevant communities of practice and will share information with the AIGB as appropriate.
- The NRC currently retains access to industry expertise across a variety of technical domains including AI. License holders will periodically obtain and share information with the AIGB as appropriate to inform decision-making and awareness.
- The NRC is working with the General Services Administration (GSA) Center of Excellence to perform an AI maturity assessment in accordance with the requirements of M-24-10. The findings, observations, and recommendations resulting from this effort will provide input into the agency's AI strategic planning.

|  RETURN TO LIST [\[#scope\]](#)

AI Use Case Inventories

NRC process for soliciting and collecting AI use cases across all sub-agencies, components, or bureaus for the inventory:

- **Idea Collection:** The NRC leverages an application named Ideascale for the capture of staff ideas pertaining to the potential use of AI. These ideas are periodically assessed by the AI Use Case Review Team in consultation with relevant functional subject matter experts to determine the feasibility of the ideas and to make a determination as to whether the idea will be elevated to or included within an agency AI use case.
- **IT Service Management:** The NRC's IT Service Managers and IT Product Owners may periodically identify opportunities to leverage AI to support IT service operations and continual service improvement. These use cases will be included within the use case inventory as appropriate.
- **Focus Groups:** The NRC will engage specific functional areas to identify opportunities to leverage AI in areas where the agency is building technical capacity. This will allow use cases to be focused on maximizing the use of emerging technologies for a specific subset of business needs. Use cases identified through this engagement will be included within the use case inventory as appropriate.
- **Program Office Requests** – The NRC will capture and process requests for the delivery of AI capabilities via the agency's digital governance process (known as Intake). This process engages IT policy subject matter

experts from a variety of domains in the review of all IT service requests that may result in a change to the agency's architecture.

To ensure that the NRC inventory is comprehensive, complete, and encompasses updates to existing use cases, the NRC will:

- Provide a single process for capturing and communicating use cases to the CAIO and AIGB as appropriate.
- Maintain detailed use case metadata as required.

|  RETURN TO LIST [\[#scope\]](#)

Reporting on AI Use Cases Not Subject to Inventory

NRC process for soliciting and collecting AI use cases that meet the criteria for exclusion from being individually inventoried, as required by Section 3(a)(v) of M-24-10:

- While the NRC aims to maintain a transparent inventory of AI use cases, certain use cases may be excluded based on any one or more of the following:
 - **Mission Risk:** Use cases that, if disclosed, could negatively impact or create risks to the agency's mission, employees, customers, or the public.
 - **Confidentiality Agreements:** Use cases subject to confidentiality agreements with other agencies, customers, employees, or stakeholders.
 - **Security Concerns:** Use cases that involve sensitive or classified information that cannot be publicly disclosed.

NRC plans to periodically revisit and validate AI use cases:

- **AIGB Reviews:** The NRC will conduct semi-annual reviews of the AI use case inventory to coincide with AIGB meetings to identify any changes or updates needed. The AIGB may also be convened outside of the standard schedule should the need arise.
- **Validation Criteria:** The NRC will leverage predefined criteria (as described above) to reassess use cases and determine whether previously excluded cases should be included or whether any new cases meet the exclusion criteria.

|  RETURN TO LIST [\[#scope\]](#)

Advancing Responsible AI Innovation

AI Strategy

The NRC is working with the General Services Administration (GSA) Center of Excellence to perform an AI maturity assessment in accordance with the requirements of M-24-10. The findings, observations, and recommendations resulting from this effort will provide input into the agency's AI strategic planning.

|  RETURN TO LIST - [\[#scope\]](#)

Removing Barriers to the Responsible Use of AI

Barriers to the responsible use of AI identified by the NRC to include steps the NRC has taken (or plans to take) to mitigate or remove these identified barriers:

- **Governance** – The NRC must continue to identify and mitigate the risks associated with the use of AI capabilities. To address this, the agency must acquire dedicated resources with adequate expertise to understand the technology and its risks in the context of the potential impact to the agency's mission and its stakeholders. The NRC will continue to communicate its budgetary needs in an attempt to fill resource gaps in governance development.
- **Digital Expertise** – The NRC must build an AI workforce that can manage AI capabilities as a service. To address this, the agency must acquire dedicated resources with adequate expertise to understand the technology and how it can be responsibly and ethically applied to the varied business functions of the agency. These resources must also be able to manage the lifecycle of these capabilities while modernizing them at the pace of industry. The NRC will continue to communicate its budgetary needs in an attempt to fill resource gaps in AI application development and service management.
- **AI-Enabling the Workforce** – Though the NRC's workforce has expressed interest in AI capabilities, it has also expressed trepidation as well as a general lack of knowledge of AI capabilities. To address this, the agency must continue to enable effective change management to enable the workforce to take full advantage of AI capabilities as they are introduced.
- **Develop an AI Ecosystem** – The NRC must develop a data and application infrastructure to support AI development, testing, and implementation. To address this, the agency will continue to identify existing and needed resources (e.g., tools, libraries, monitoring capabilities, etc.) while communicating the associated budgetary requirements via budget formulation and execution processes.
- **Government Availability** – Many AI capabilities that are available to the public and/or the commercial sector are not available to the government sector at the same time due to the more stringent cybersecurity

requirements. There will always be a lag between the availability of these capabilities in the market and the NRC's use of these capabilities.

- **Budget** – The NRC is only able to assess, test, implement, and maintain new capabilities where resources have been made available to do so. Budgetary constraints and resource prioritization associated with the delivery of IT/IM services limit the agency's capacity for innovation and modernization. The CIO/CAIO will continue to express the need for budgetary resources during budget formulation and execution cycles.

Status of internal guidance for the use of generative AI:

- The NRC has developed internal guidance for the use of generative AI in the form of AI Rules of Behavior which is accompanied by a training course designed to inform staff of the capabilities and risks of generative AI. These rules of behavior will be incorporated into the agencywide IT rules of behavior.

|  RETURN TO LIST [\[#scope\]](#)

AI Talent

Planned or in-progress to increase NRC AI talent:

- The NRC is leveraging the Direct Hire Authority to acquire additional staff to support AI activities. Additional FTE in this area were acquired via the FY25 OMB passback and have been requested in the FY26 budget.
- The agency seeks to expand its AI workforce with expertise in Business Analysis, AI-enabled Knowledge Discovery, AI-enabled Content Generation, AI Application Testing, Development, and Monitoring.
- The agency's AI Strategic Leadership Council is responsible for assessing the agency's resource needs and is co-chaired by the Office of the Chief Information Officer and the Office of Nuclear Regulatory Research.
- NRC has developed an AI Competency Model in accordance with the AI Competency Model developed by the Office of Personnel Management (OPM). The agency will use the competency model as an input into its efforts to upskill and reskill existing and new staff.

NRC's plans to provide resources or training to develop AI talent internally and increase AI training opportunities for Federal employees:

- The NRC is in the process of expanding its learning management toolset via the acquisition of a new learning management platform. In combination with existing capabilities, this new platform will expand learning opportunities for management and staff in the use of AI capabilities.

- Via the CAIO Council, the agency is taking advantage of leadership training in AI that has been made available to the federal government via conferences, online courses, instructor-led classes, and symposiums.
- The NRC is currently leveraging training support staff to provide agencywide brown bag learning sessions on a variety of AI capabilities.

|  RETURN TO LIST - [\[#scope\]](#)

AI Sharing and Collaboration

NRC's position on for ensuring that custom-developed AI code—including models and model weights—for AI applications in active use is shared consistent with Section 4(d) of M-24-10:

- At this time, the NRC has no intentions of developing custom-developed AI code. The goal is to leverage commercial off the shelf (COTS) AI tools and services. Though unlikely, if the agency does venture into custom-developed AI code, the agency's governance processes will include provisions to require that code is shared consistent with the Section 4(d) of M-24-10.

NRC's position on encouraging or incentivize the sharing of code, models, and data with the public:

- At this time, the NRC has no intentions of developing custom-developed AI code. The goal is to leverage commercial off the shelf (COTS) AI tools and services. Though unlikely, if the agency does venture into custom-developed AI code, the agency's governance processes will include provisions to require that code, models and data are shared consistent with the Section 4(d) of M-24-10.

|  RETURN TO LIST - [\[#scope\]](#)

Harmonization of Artificial Intelligence Requirements

NRC actions taken to document and share best practices regarding AI governance, innovation, or risk management:

- The NRC will leverage the AIGB, the AI Steering Committee, and the AI Strategic Leadership Council to share best practices. As they are developed, these best practices will be embedded within their respective domains to ensure that they can be incorporated into the agency's relevant standard operating procedures.

|  RETURN TO LIST - [\[#scope\]](#)

Managing Risks from the use of Artificial Intelligence

Determining Which Artificial Intelligence Is Presumed to Be Safety-Impacting or Rights-Impacting

NRC process for determining which AI use cases are rights-impacting or safety-impacting:

- The NRC assesses AI use cases based on a combination of the type of AI to be applied, the data to be leveraged, and the desired outcomes of the use case. It is the combination of these factors, compared against the purposes that are presumed to be safety-impacting and the purposes that are presumed to be rights-impacting that illustrates the potential impact to human safety and/or human rights.

NRC-specific criteria to guide a decision to waive one or more of the minimum risk management practices for a particular use case:

- The NRC has not developed distinct criteria to guide a decision to waive one or more of the minimum risk management practices as the agency currently has no need nor intention of providing opportunities for waivers.

NRC process for issuing, denying, revoking, tracking, and certifying waivers for one or more of the minimum risk management practices:

- Per the statement above, this is not applicable to the NRC.

|  [RETURN TO LIST](#) [\[#scope\]](#)

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

NRC controls to prevent non-compliant safety-impacting or rights-impacting AI from being deployed to the public:

- The NRC currently does not anticipate deploying any AI-enabled technology to the public.
- The NRC will leverage its process for determining which AI is presumed to be safety-impacting or rights-impacting prior to the implementation of any AI capability.

NRC's intended process to terminate, and effectuate that termination of, any non-compliant AI:

- Should an AI tool or service be deemed non-compliant, the agency’s governance process will ensure that it will be removed from the production environment following engagement and collaboration with impacted stakeholders.
- In accordance with the agency’s developing governance process, the NRC will assess the conditions in which the tool or service entered into a state of non-compliance and work to identify a remediation plan. If non-compliance cannot be remediated, an alternate approach will be sought to deliver the desired outcomes.

|  RETURN TO LIST - [\[#scope\]](#)

Minimum Risk Management Practices

NRC’s plans to document and validate implementation of the minimum risk management practices:

- The NRC will manage the use of AI capabilities in accordance with its existing practices for IT Service Management. All AI tools and services will be overseen by IT Service Manager(s) with responsibility for Service Operation and Continual Service Improvement which will encompass the minimum risk management practices.
- In alignment with existing internal IT Service and IT Investment reporting practices, the IT Service Manager(s) and their respective team(s) will provide periodic updates to the CIO and CAIO while raising topics for AIGB consideration.

|  RETURN TO LIST - [\[#scope\]](#)

SECTION NAVIGATION +



Choose a section

Page Last Reviewed/Updated Thursday, September 19, 2024