



GSA Order: Use of Artificial Intelligence at GSA

CIO 2185.1A
GSA-IT
caio@gsa.gov

Purpose

This directive establishes the governing policies regarding the controlled access and responsible use of artificial intelligence (AI) technologies and platforms. It addresses the assessment, procurement, usage, monitoring, and governance of AI systems and software within the GSA network, in conjunction with all existing security, privacy, policies, directives, ethics regulations, and laws.

Background

The AI in Government Act of 2020 (Public Law 116-260), AI Training Act of 2023 (Public Law 117-207), Executive Order 13859, Executive Order 13960, Executive Order 14110, Executive Order 14091, M-21-06, M-24-10, OMB Circular No. A-119, and the AI Bill of Rights direct all Federal agencies to:

1. Ensure that all AI and automated systems comply with applicable Federal law in a manner that advances equity, safety, and privacy;
2. Establish or update processes to measure, monitor, evaluate, and report on AI activities, use-cases, their ongoing performance, and manage the risks of using AI through regular risk assessments as required, especially for safety-impacting and rights-impacting AI;
3. Prioritize appropriate uses of AI that improve their agency's mission, advance equity and identify and remove barriers to the responsible use of AI in the agency, including through the advancement of AI-enabling enterprise infrastructure, workforce development measures, policy, and other resources for AI innovation;
4. Ensure adequate infrastructure and capacity is available to sufficiently curate agency datasets for AI usage, including the requisite data governance and management practices as they relate to data curation, labeling, and stewardship;

5. Initiate measures and procedures to regularly assess the agency's AI workforce capacities and its projected AI workforce needs;
6. Support interagency coordination bodies related to AI activities and AI standards-setting initiatives, and encourage agency adoption of voluntary consensus standards for AI.

Applicability

This order applies to:

1. All GSA employees and contractors that may have a need to access or share data, as well as system-to-system data exchanges;
2. IT systems owned and operated by or on the behalf of any of the GSA Service and Staff Offices (SSOs), including Regional Offices;
3. GSA or Federal data contained on or processed by IT systems owned and operated by or on the behalf of any of the GSA SSOs, including Regional Offices;
4. The Office of Inspector General (OIG) to the extent that the OIG determines it is consistent with the OIG's independent authority under the Inspector General Act of 1978 (5 U.S.C. App. 3) and does not conflict with other OIG policies or the OIG mission; and
5. The Civilian Board of Contract Appeals (CBCA) only to the extent that it is consistent with the CBCA's requisite independence as defined by the Contract Disputes Act (CDA) and its legislative history. 41 U.S.C. §§ 7101-7109 (2012) and S. Rep. No. 95-1118 (1978).

Cancellation

This directive cancels the Security Policy for Generative Artificial Intelligence (AI) Large Language Models (LLMs) (Number: CIO IL-23-01).

Roles and Responsibilities

1. Chief AI Officer (CAIO): in addition to the responsibilities defined in Section 8(c) of EO 13960 and Section 4(b) of EO 14091, the CAIO must:
 - a. Maintain awareness of AI activities within GSA, including how the systems work, how they were designed, and what specific purposes they serve;
 - b. Establish and update processes to measure, monitor, and evaluate the performance, accessibility, equity, cost, and outcomes of AI applications;

- c. Establish, maintain, and chair AI oversight governing bodies;
 - d. Issue AI compliance plans and oversee agency compliance with the AI Executive Order 14110;
 - e. Oversee the development of GSA's AI inventory and other necessary reporting; and
 - f. Identify and convene external individuals or organizations with AI expertise who can provide expert input to agency officials that is relevant to GSA mission functions.
 - g. Issue waivers for individual applications of AI, in coordination with other officials responsible for those AI applications, from elements of Section 5 of M-24-10.
 - h. Establish, and maintain over time, criteria for categories of individual applications of AI that do not require disposition through the AI Governance Board or AI Safety Team.
2. [AI Governance Board](#): co-chaired by the CAIO and the Deputy Administrator of GSA, shall include representation from senior agency officials responsible for key enablers of AI adoption and risk management, including at least IT, cybersecurity, data, human capital, legal, procurement, budget, agency management, customer experience, performance evaluation, statistics, risk management, equity, privacy, civil rights and civil liberties, the Office of the Inspector General, and officials responsible for implementing AI within an agency's program office. The AI Governance Board Charter will define the roles and responsibilities of the AI Governance Board.
3. [AI Safety Team](#): The working group reporting to the Chief AI Officer in their role as co-chair of the AI Governance board, responsible for adjudicating use cases, developing draft guidance, policy, and standards. The AI Safety Team will comply with existing Federal and agency Security and Privacy policies when making use case dispositions. The AI Safety Team shall be populated by delegated representatives of the AI Governance Board and the CAIO.
- a. The AI Safety team will be composed of individuals who can provide diverse perspectives on the use of AI, including developers, architects, data scientists, user experience/customer experience experts, privacy, security and both internal and public mission staff.

- b. The AI Safety Team will be empowered to independently adjudicate Familiarization, Pre-Acquisition, and Research and Development use cases.
 - i. The AI Safety team is responsible for providing disposition recommendations for Production or Production-Intent use cases.
 - ii. All determined Rights or Safety Impacting use cases must be ultimately adjudicated by the AI Governance Board.
 - c. The AI Safety Team shall enforce all GSA-authorized security, privacy, and audit policies to protect CUI and ensure GSA IT systems operate within acceptable levels of residual risk. These include, but are not limited to:
 - i. Privacy Threshold Assessments (PTAs);
 - ii. Privacy Impact Assessments (PIAs);
 - iii. Privacy Act Statements;
 - iv. System of Records Notices (SORNs);
 - v. Authorizations to Operate (ATOs); and
 - vi. FedRAMP authorizations;
 - d. The CAIO must review and approve all production or production-intent use cases. The CAIO and the AI Governance Board maintain full access to all use cases registered with the AI Safety Team and can review any use case at any point.
 - e. AI use cases deemed to have significant implications for rights or safety by the AI Safety Team, the CAIO, or the AI Governance Board will be adjudicated as Rights-Impacting or Safety-Impacting, and the AI use will be subject to additional monitoring, reporting, and review processes.
4. System Owner: Shall be responsible for reporting all AI use cases for review by the AI Safety Team and providing updates should any significant modifications to the AI system occur or if the AI system is decommissioned.
5. Executive Sponsor: Shall be named sponsors for AI use cases, and ensure alignment with the strategic objectives, risk posture, and resourcing priorities of the AI Governance Board. Executive Sponsors are not required for familiarization use cases.

6. Authorized Users of IT Resources:

- a. General Practitioner: Shall be responsible for protecting federal nonpublic information, reporting any potential IT security incident, adhering to [GSA's Information Technology \(IT\) General Rules of Behavior](#) and all provisions in this directive. All AI users also have a responsibility to report any use of AI to the AI Safety Team if they believe the use case has not already been registered by the System Owner.
- b. Specialized Practitioner: In addition to all responsibilities of a general AI practitioner, a specialized practitioner shall be responsible for implementing and maintaining all GSA IT software development and security standards when supporting the development and implementation processes of AI software and solutions.

Signature

/S/_____

David Shive
Chief Information Officer
Office of GSA IT

6/7/2024_____

Date

THIS PAGE IS INTENTIONALLY LEFT BLANK

Table of Contents

Purpose	1
Background	1
Applicability	2
This order applies to:	2
Cancellation	2
Roles and Responsibilities	2
Signature	5
Table of Contents	7
1. Introduction	9
1.1 Objectives	9
1.2 Scope	9
1.3 Principles	10
2. Policy	10
2.1 General AI usage	10
2.2 New or Proposed AI Use Cases	12
2.3 Existing AI Use Cases	13
2.4 AI Code and Models	14
2.5 Data assets and sources	15
2.5.1 Internal Data Assets	15
2.5.2 External Data Assets	16
2.5.3 AI-Generated Data Products	16
2.5.4 Data Dissemination Requirements	17
2.6 Responsible Procurement of AI	17
2.6.1. Pre-Acquisition	17
2.6.2. Procuring AI	17
2.6.3 Procurement policy updates	17
2.7 Tool or Product AI Enhancements	17
2.8 Publication Requirements	18
2.9 Minimum Requirements for Either Safety-Impacting or Rights-Impacting AI	18
2.9.1 Additional Requirements for Rights-Impacting AI	20
2.9.2 Excepted scenarios for Rights-Impacting or Safety-Impacting AI use cases	20
2.9.3 Use-Case Waivers	20
2.10 Organizational Risk Tolerance and Use Case Risk Rubric	21
3. Legal and Programmatic Authorities	21
4. Definitions	22
5. Appendix A: Presumed Rights-Impacting and Safety Impacting Use Cases	27
5.1 Rights-Impacting Use Case Examples	27
5.2 Safety-Impacting Use Case Examples	29

6. Appendix B: AI Impact Statement Guidance

30

7. Appendix C: Additional Documents

31

1. Introduction

As Artificial Intelligence (AI) technologies continue to evolve and expand into the workflows of the federal government, it is crucial that the use of these technologies are managed to maximize effectiveness while minimizing potential harm and managing or mitigating potential risks. AI has the potential to augment or improve mission delivery, service offerings, and productivity across all GSA equities. However, without oversight, controls, and human intervention protocols in place, AI can also cause harm by introducing or reinforcing discriminatory practices, invading people's privacy, or enabling disinformation to propagate at scale. To mitigate these potential issues safely while capitalizing on the potential benefits of these emerging technologies, policy controls must be established for the safe, secure, equitable, and trustworthy development and use of AI.

This directive outlines the controls for AI usage within GSA, the governance and oversight infrastructure required to enable the responsible use of AI, the processes available to GSA employees in developing AI use cases for mission work, and the disclosure requirements for all AI implementations.

1.1 Objectives

The objectives of this order are to:

- a. Define AI governance model and procedures necessary to promote the safe, equitable, and responsible use of AI technologies while managing its associated risks for GSA business activities;
- a. Enable use of AI that improves service delivery and public trust in government;
- b. Establish the roles, responsibilities, and reporting structures of the requisite oversight and governing groups;
- c. Outline the requirements of all AI systems, with noted focus on rights-impacting and safety-impacting AI systems; and
- d. Define core AI terms and concepts.

1.2 Scope

This order provides guidance for the program operations of GSA that have direct or indirect responsibility for or control over any action, activity or program that relates to AI systems, including the procurement, management, or development activities. This policy is designed to work with existing IT security and privacy policies.

1.3 Principles

This directive is based on the principles of public trust, scientific integrity, risk management, equity, transparency, safety, and collaboration. AI systems must be developed and deployed in a manner that prioritizes the public good while also taking into account the potential risks and benefits. These principles are essential to ensure the safe, responsible, and ethical development and deployment of AI systems across GSA.

2. Policy

AI use cases in GSA are categorized as follows:

1. **Familiarization:** working with AI for professional development and training using non-sensitive, public information and publicly available tools. These use cases are relegated for individual uses, with the specific goal of gaining familiarity with market offerings, and are most closely aligned with professional training activities;
2. **Pre-acquisition activity:** assessing or piloting the capabilities of an AI system or performing market analyses before acquiring the technology. This includes Request for Information (RFIs), industry days, or any scenario where third party developers or vendors provide demonstration products outside of GSA's network or infrastructure. These use cases can not use non-public Federal Controlled Unclassified Information (CUI) data or interface with internal GSA systems;
3. **Research and Development:** work involving the development of a capability using internal systems, processes, and data, but without the immediate intent to promote the research output to a production environment or workflow. These use cases can not support GSA business activities directly and may only take place in the Enterprise Data Solution environment or approved research environments; and
4. **Production or production-intent:** use cases involving the incorporation of AI for deployment into production environments or workflows. The work products of these use cases directly support GSA business activities.

2.1 General AI usage

For all use cases, individuals acting on behalf of GSA must register every proposed use case via GSA's AI Request Form. Use case requests are assessed by the AI Safety Team, which identifies each use case's risk profile and adjudicates use cases classified

as Familiarization, Pre-Acquisition, Research and Development, and Production or Production-Intent. The AI Safety Team may request additional guidance from the Chief AI Officer (CAIO) and AI Governance Board as necessary.

- a. General access to publicly available, approved third-party AI endpoints and tools shall be blocked from the GSA network and GFE devices.
 - i. Access will be made available upon completion of GSA's AI request form, detailing intended usage and acknowledgment of the requirements of this Directive and GSA's IT General Rules of Behavior.
 - ii. Only endpoints and tools approved by the CAIO, the CISO, and the [AI Governance Board](#) will be made available.
 - iii. Access to public interfaces will only be approved for uses that involve publicly available data and are for familiarization purposes only. No output from publicly available products or tools may be introduced as a GSA production work product without approval from the AI Governance Board.
 - iv. Federal nonpublic information (including work products, emails, photos, videos, audio, and conversations that are meant to be pre-decisional or internal to GSA), such as controlled unclassified information (CUI), personally identifiable information (PII), and Business Identifiable Information (BII), shall not be used as inputs (e.g. prompts or training data) to any AI system without prior authorization from the [AI Governance Board](#).
- b. Any work product outputs materially modified by or solely produced by generative AI systems must be labeled or watermarked in a manner that makes the recipient aware of the system(s) involved and whether they edited or authored the work. Content types include:
 - i. Data;
 - ii. Code;
 - iii. Text (e.g. temporary and permanent records);
 - iv. Applications (e.g. chatbots, recommendation engines, etc.);
 - v. Audio;
 - vi. Imagery; and

- vii. Video.
- c. All production systems using AI capabilities that provide direct interface with the public must include:
 - i. Notice and explanation of its services written in [plain language](#); and
 - ii. Human alternatives or fallback options where practicable.
- d. All AI software must have a valid Authorization to Operate prior to use for Research and Development and Production use cases.
- e. Any output from LLMs used to generate code or content to be published on federal internet or intranet pages or to be used in Agency Official Communications (as defined in 36 CFR 1194 E205.3), shall be manually reviewed to ensure that code or the content conforms to Section 508 of the Rehabilitation Act of 1973 and to the Section 508 Technical Standards for ICT Accessibility.

2.2 New or Proposed AI Use Cases

All AI use case requests must be submitted to the AI Safety Team via the AI Request Form. If a model that is not currently authorized is being requested, the use case must also submit an AI Model Request Form. Research and Development use case requests shall also submit an [Experimental Design Statement](#). All applicants must include in their submission the following information:

- a. Category of use case type (i.e., familiarization, pre-acquisition, research and development, or production); and
- b. Intended purpose for the AI and the expected benefit;
- c. The creator of the AI system;
- d. The environment(s) the AI system will be located in;

Pre-acquisition, research and development, and production AI use cases are required to provide additional information, including:

- e. What specific metrics or qualitative measures will be used to assess impact, employing performance measurement or program evaluation methods;
- f. Intended user/audience of the AI system or AI capability;

- g. Justification for how the AI is better suited to accomplish the relevant task than alternative methods;
- h. What risks are associated with the use of an AI in the requested use-case and what measures should be employed to reduce or mitigate the risks; and
- i. What data will be used by the AI in the use case.

The AI Request Form may be modified to require additional or different information as deemed necessary by the AI Governance Board.

2.3 Existing AI Use Cases

Every year, all existing use cases are required to re-register with the AI Safety Team via the AI Request Form, with the exception of familiarization use cases. Use cases must also submit the same form to report ceasing operations. The specific requirements are as follows:

- a. All existing AI use cases, with the exception of familiarization use cases, shall be reported to the [AI Governance Board](#) on an annual basis.
- b. Any use case that undergoes a significant modification must be re-submitted to the AI Governance board for reassessment.
- c. Any use case where there has been a cybersecurity or privacy incident must be re-submitted to the AI Governance board for reassessment within 30 days of the reported incident;
- d. AI systems that use nonpublic information shall be conducted within approved secure enterprise systems, such as the Enterprise Data Solution (EDS).
- e. All AI systems are subject to independent system reviews and assessments of the use case, the system and its architecture, the security protocols, and privacy measures upon request by the:
 - i. CAIO;
 - ii. AI Governance Group or designee;
 - iii. Chief Information Security Officer;
 - iv. Chief Technology Officer; and
 - v. Chief Privacy Officer.

2.4 AI Code and Models

All internally developed AI code shall be shared for internal consumption as well as open sourced in public repositories. All code shall adhere to GSA's Open Source Software (OSS) Policy (2107.1 CIO) before sharing code.

- a. All custom-developed code - including models and model weights - for AI applications shall be:
 - i. shared internally; and
 - ii. open-sourced to the public;
- b. Code and models no longer in active use may be archived and do not need to be maintained.
- c. Use of Open Source or COTS models:
 - i. Shall be approved for use by the AI Safety Team through link to AI Model Request Form;
 - ii. Shall be treated as a system integration; and
 - iii. Shall adhere to standard architectural protocols and requirements.
- d. Exceptions to this provision include:
 - i. the sharing of code is restricted by law or regulation;
 - ii. sharing of code would create an identifiable risk to national security, confidentiality of Government information, individual privacy, or the rights or safety of the public;
 - iii. the code or models were used for Research and Development use cases;
 - iv. contractual obligations that prevent the sharing of code; and
 - v. the sharing of code would create an identifiable risk to agency mission, programs, or operations, or the stability, security, or integrity of an agency's system or personnel.

2.5 Data assets and sources

System owners of research and development or production AI systems shall report on the data used in the design, development, training, testing, and operation of an AI system. This includes:

- a. what data are being used;
- b. the purpose of the data being used within the AI system;
- c. who are the owners of the data being used;
- d. how the data are relevant to the task being automated; and
- e. what is the sensitivity level of the data required for the AI use case.

The AI Safety Team or CAIO may request additional information about the data being used, including:

- a. documentation on the data collection and preparation process, including the data provenance;
- b. measures of the quality and representativeness of the data for its intended purpose;
- c. how the data will be used for the AI's development, testing, and operation; and
- d. measures that demonstrate that the data adequately cover real-world scenarios, and how shortcomings are being addressed.

All data used in Research and Development or Production use cases for an AI system's design, development, training, testing, and operation shall be:

- a. registered and published in the EDS catalog; and
- b. adhere to the [Internal Data Sharing Directive](#) and its data categorization framework.

2.5.1 Internal Data Assets

All internal data assets are subject to the following specific requirements:

- a. no internal data assets may be used as input for public AI systems; and

- b. no sensitive data (e.g. PII, CUI, Procurement Sensitive) may be used with AI systems without clearance from the AI Safety Team, a submission of an AI impact statement, and a valid ATO.

2.5.2 External Data Assets

For data sources proposed to be used by an AI system that are generated from external sources (i.e., non-GSA owned and maintained data assets), the AI system manager shall report on the following:

- a. The originator, collection methodology, and preparation process of all external data sources shall be registered with the AI Safety Team. These specifics shall be resubmitted:
 - i. during the annual resubmission process to continue usage; and
 - ii. if any significant modification occurs to the use case.
- b. The data sources shall be maintained, indexed, and made available via the Enterprise Data Solution.

2.5.3 AI-Generated Data Products

All AI-generated data outputs or products must be labeled as such in its metadata, and indexed and cataloged in the EDS system for internal discovery purposes. This includes any generated modification to existing data products. Datasets that have undergone augmentation from an AI system, such as data imputation or field creation and population, must include notice in the metadata holdings as to which records were modified or created, and by what system, including the AI systems version information. All AI-generated data must adhere to existing data, privacy, and security policies.

Exceptions to AI-generated content notice may include, but are not limited to:

1. Metadata, including field titles, descriptions, and domain associations;
2. Classification or tagging labels for discovery or findability purposes; and
3. Domain association for general data ontology management purposes.
4. Data authored by humans or non-AI systems which may contain generated content that does not fundamentally challenge the authorship of the data, such as emails or chats which contain auto-completed text.

2.5.4 Data Dissemination Requirements

Data used in the development of AI models or applications shall be qualified as a data asset under the definition of the Open, Public, Electronic, and Necessary (OPEN) Government Act, and shall be publicly released as an open government data asset on data.gov. Exceptions for this provision would follow the safety and security considerations in Section 4.7 of EO 14110. All existing risk mitigation and privacy process controls remain in force for all data products identified for dissemination.

2.6 Responsible Procurement of AI

2.6.1. Pre-Acquisition

For a GSA-funded procurement, market research should be used to determine if AI will be offered as a solution or potential solution to the planned procurement. If it is determined during market research that AI may be proposed by an offeror as part of their total solution, acquisition teams must coordinate the acquisition plan and solicitation with the CAIO.

Any procurement considerations regarding AI usage at GSA must be submitted to the CAIO and reviewed by the AI Safety Team before proceeding. In accordance with General Services Acquisition Manual (GSAM) 507.104(a)(6), acquisition plans contemplating the procurement of AI for use at GSA must be coordinated and approved by the CAIO. All plans shall be submitted to the AI Safety Team.

2.6.2. Procuring AI

In accordance with GSAM 507.104(a)(6) and 511.170, prior to release of a solicitation for AI for use at GSA, the acquisition team must ensure the requirements document (Performance work Statement (PWS)/Statement of Objective (SOO)/Statement of Work (SOW)) has been coordinated and approved by the CAIO. Submit solicitations to the AI Safety Team. A solicitation can not be released until the CAIO has provided written approval.

2.6.3 Procurement policy updates

The procurement policy for AI will be updated over time in alignment with GSA guidelines and directives, and will maintain compliance with federal acquisition standards.

2.7 Tool or Product AI Enhancements

In many cases, products or tools that have already been procured and have an active Authority to Operate (ATO) will be enhanced with AI.

- a. All existing tools with ATO that receive AI enhancements must:
 - i. be submitted by the System Owner to the Authorizing Official for assessment and receive a reauthorization prior to bringing the new functionality into the ATO boundary; and
 - ii. report the AI capability as a procurement via the AI Request Form.
- b. The application of the AI enhancement will be reviewed and dispositioned by the AI Safety Team.
- c. Should the AI enhancement violate policy outlined in this directive, the enhancement will be required to be turned off or the software be reverted to version that does not contain the enhancement. If it is not possible to revert the AI enhancement, process controls and policy will be required to be submitted by the system owner to the CAIO, proving the enhancement is not used in the applicable use cases.

2.8 Publication Requirements

All research and development and production or production-intent AI systems currently in use will be publicly disclosed pursuant to Section 3(a) of [M-24-10, "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence."](#) These use cases will be included in the AI use case inventory and hosted on gsa.gov. All use cases shall provide the data elements required by OMB and its Integrated Data Collection process or any OMB-designated superseding processes.

Exceptions for publication include AI systems whose disclosure would be inconsistent with applicable law and governmentwide policy.

Aggregated statistics of all use cases will be disclosed publicly, to include but not limited to:

- a. the number of rights-impacting and safety-impacting use cases currently in operation;
- b. the compliance status of all use cases; and
- c. all use case waivers currently in force.

2.9 Minimum Requirements for Either Safety-Impacting or Rights-Impacting AI

All AI use cases that match the definitions of “safety-impacting AI” or “rights-impacting AI” are subject to the additional requirements in this section because of the potential risk they can pose, for example, discrimination and other harms to people. Recognizing both the risks and opportunities presented by potential “safety-impacting AI” or “rights-impacting AI” capabilities, the CAIO is establishing transparent governance and compliance processes that responsibly address the full scope of these potential risks. System owners and their designees are responsible for enacting these minimum requirements. [Appendix A](#) identifies use cases that would be presumed as covered AI (e.g. either rights-impacting or safety-impacting).

Waivers from minimum practices may be requested. All requests must be made to the CAIO and AI Governing Board who will adjudicate the request. Any covered AI not in compliance by December 1, 2024 shall cease operations until compliant with the following controls.

All rights-impacting and safety-impacting use cases must follow these practices *before* employing the AI into any use case:

- a. Complete an AI [Impact Statement](#);
- b. Submit an AI system test plan that demonstrates real-world context testing, and contestability as necessary;
- c. Submit to an independent evaluation of the AI system from the CAIO or their designee;

All rights-impacting and safety-impacting use cases must follow these practices *while* employing the AI into any use case:

- d. Conduct ongoing monitoring of the AI system and establish thresholds for periodic human review;
- e. Mitigate emergent risks to rights and safety identified through routine testing, continuous monitoring protocols, or third-party findings;
- f. Ensure all system practitioners have taken requisite AI training requirements;
- g. Include human validation and intervention protocols to ensure all output decisions made by AI systems are regularly evaluated by system practitioners; and

- h. Provide public notice and plain language documentation regarding the rights- or safety-impacting use case through the public interface, in public disclosure statements, and the AI use case inventory.

2.9.1 Additional Requirements for Rights-Impacting AI

AI use cases deemed as Rights-Impacting must follow these additional requirements *before* implementation:

- a. Proactively identify and mitigate algorithmic discrimination or bias;
- b. Assess and mitigate disparate impacts for protected classes;
- c. Conduct direct user testing of system interactions; and
- d. Solicit comments from the user community and conduct post-transaction customer feedback activities in coordination with the Office of Customer Experience or equivalent.

AI use cases deemed as Rights-Impacting must follow these additional requirements *while* employing the AI into any use case:

- e. Conduct ongoing monitoring studies for AI-enabled discrimination;
- f. Notify any negatively affected individuals;
- g. Provide fallback and escalation options for AI processes or outcomes; and
- h. Provide opt-out alternatives where practicable.

2.9.2 Excepted scenarios for Rights-Impacting or Safety-Impacting AI use cases

The following Rights-Impacting or Safety-Impacting AI use cases do not need to follow the requirements set out in 2.9 and 2.9.1 above:

- a. Evaluation of a potential vendor, commercial capability, or freely available AI capability that is not otherwise used in agency operations, solely for the purpose of making a procurement or acquisition decision;
- b. Evaluation of a particular AI application because the AI provider is the target or potential target of a regulatory enforcement action; and
- c. Research and development purposes.

2.9.3 Use-Case Waivers

- a. The CAIO may waive one or more of the stated requirements for specific covered AI applications with conditions in scenarios where one or more of the requirements would increase risks to safety or rights overall or would create an unacceptable impediment to critical agency operations.
 - i. Appeals for waivers may be submitted by System Owners or delegates with written justifications to the CAIO.
 - ii. All waivers must be centrally tracked and are subject to publication requirements outlined in [Publication Requirements](#).
 - iii. All waivers will be reassessed on an annual basis.

2.10 Organizational Risk Tolerance and Use Case Risk Rubric

The AI Governance Board shall establish the enterprise's AI risk tolerance, prioritization, and risk management strategic approach. All risk management activities shall comport with [Enterprise Risk and Strategic Initiatives \(ERSI\) Board](#) reporting requirements, under [GSA's Enterprise Risk Management \(ERM\) Policy](#). This includes:

- a. Establishing likelihood and impact ranking criteria and thresholds;
- b. Defining the considered factors for the use case risk rubric; and
- c. Establishing the risk management practices and processes that are required for AI systems;

The AI Safety team is responsible for assessing use cases based on the guidance provided by the AI Governance Board.

Each System Owner is responsible for implementing the risk management processes defined by the AI Governance Board.

3. Legal and Programmatic Authorities

- a. Executive Order 14110. Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. October 30, 2023.
- b. Executive Order 14091. Further Advancing Racial Equity and Support for Underserved Communities Through the Federal Government. February 2023.
- c. Executive Order 13960. Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government. December 2020.

- d. Executive Order 13859. Maintaining American Leadership in AI. February 2019.
- e. The AI in Government Act of 2020 (Public Law 116-260).
- f. AI Training Act of 2023 (Public Law 117–207).
- g. Generative AI and Specialized Computing Infrastructure Acquisition Resource Guide.

4. Definitions

1. AI Use Case: The application of artificial intelligence technology to address specific challenges or improve existing processes within the agency. This can include automating repetitive tasks, improving data analysis and decision-making, and enhancing customer service through chatbots or virtual assistants. Examples of AI use cases within GSA could include using machine learning algorithms to optimize procurement processes, or leveraging natural language processing to improve search functionality on the agency's website.
2. System Owner: System Owners are GSA management officials with responsibility for the acquisition, development, maintenance, implementation, and operation of GSA's IT systems. System Owners cannot be Information System Security Officers (ISSOs) or Information System Security Managers (ISSMs). System Owners represent the interests of the system throughout its lifecycle. Primary responsibility for managing risk should rest with the System Owners.
3. Algorithmic discrimination: The term “algorithmic discrimination” has the meaning established in Section 10(f) of Executive Order 14091 of February 16, 2023.
4. Artificial Intelligence (AI): The term “artificial intelligence” has the meaning established in Section 238(g) of the John S. McCain National Defense Authorization Act for Fiscal Year 2019, which states that “the term ‘artificial intelligence’ includes the following”:
 - a. Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.

- b. An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.
 - c. An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
 - d. A set of techniques, including machine learning, that is designed to approximate a cognitive task.
 - e. An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision-making, and acting.
 - f. This definition of AI encompasses, but is not limited to, the AI technical subfields of machine learning (including, but not limited to, deep learning as well as supervised, unsupervised, and semi-supervised approaches), reinforcement learning, transfer learning, and generative AI.
 - g. This definition of AI does not include robotic process automation or other systems whose behavior is defined only by human-defined rules or that learn solely by repeating an observed practice exactly as it was conducted.
 - h. For this definition, no system should be considered too simple to qualify as a covered AI system due to a lack of technical complexity (e.g. the smaller number of parameters in a model, the type of model, or the amount of data used for training purposes).
 - i. This definition includes systems that are fully autonomous, partially autonomous, and not autonomous, and it includes systems that operate both with and without human oversight.
5. Contestability: the ability to effectively challenge a decision made or augmented by AI.
6. Covered AI: AI that has been adjudicated to be Safety-Impacting or Rights-Impacting.

7. Data Asset: The term “data asset” has the meaning provided in 44 U.S.C § 3502.
8. Equity: Has the meaning established in Section 10(a) of Executive Order 14091.40
9. Federal Information: Has the meaning established in OMB Circular A-130.
10. Generative AI (GenAI): Has the meaning established in Section 3(p) of AI Executive Order 14110.
11. Production Work Product: any deliverable or tangible outcome produced as a result of work activities within a project or task. This can include documents, emails, software, presentations, reports, designs, models, and other artifacts that demonstrate progress or completion of work, measure performance, ensure quality, or facilitate communication among stakeholders. Examples of work products include, but are not limited to:
 - a. Documentation: manuals, user guides, project plans, technical specifications, meeting notes, and progress reports;
 - b. Software: code, scripts, and applications.
 - c. Designs and models: architectural blueprints, wireframes, prototypes, diagrams, and simulations;
 - d. Presentations: slide decks, infographics, dashboards, and visual aids;
 - e. Data: databases, datasets, spreadsheets, and data analysis reports;
 - f. Other deliverables: external communications, training materials, marking collateral, and audit findings.
12. Research and Development: As in OMB Circular No. A-11, Preparation Submission, and Execution of the Budget (2023), research and development is defined as creative and systematic work undertaken in order to increase the stock of knowledge—including knowledge of people, culture, and society—and to devise new applications using available knowledge.

13. Rights-Impacting AI: Has the meaning established in Section 6 of [M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence](#).

14. Risks from the Use of AI: Risks related to efficacy, safety, equity, fairness, transparency, accountability, appropriateness, or lawfulness of a decision or action resulting from the use of AI to inform, influence, decide, or execute that decision or action. This includes such risks regardless of whether:

- a. the AI merely informs the decision or action, partially automates it, or fully automates it;
- b. there is or is not human oversight for the decision or action;
- c. it is or is not easily apparent that a decision or action took place, such as when an AI application performs a background task or silently declines to take an action; or
- d. the humans involved in making the decision or action or that are affected by it are or are not aware of how or to what extent the AI influenced or automated the decision or action.

While the particular forms of these risks continue to evolve, at least the following factors can create, contribute to, or exacerbate these risks:

- a. AI outputs that are inaccurate or misleading;
- b. AI outputs that are unreliable, ineffective, or not robust;
- c. AI outputs that are discriminatory or have a discriminatory effect;
- d. AI outputs that contribute to actions or decisions resulting in harmful or unsafe outcomes, including AI outputs that lower the barrier for people to take intentional and harmful actions;
- e. AI being used for tasks to which it is poorly suited or being inappropriately repurposed in a context for which it was not intended;

- f. AI being used in a context in which affected people have a reasonable expectation that a human is or should be primarily responsible for a decision or action; and
 - g. the adversarial evasion or manipulation of AI, such as an entity purposefully inducing AI to misclassify an input.
15. Safety-Impacting AI: Has the meaning established in Section 6 of [M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence](#).
16. Significant Modification: An update to an AI application or to the conditions or context in which it is used that meaningfully alters the AI's impact on rights or safety, such as through changing its functionality, underlying structure, or performance such that prior evaluations, training, or documentation become misleading to users, overseers, or individuals affected by the system. This includes significantly changing the context, scope, or intended purpose in which the AI is used. Examples of significant modifications include, but are not limited to, changes in:
- a. production status, including but not limited to:
 - i. Any change to the use case type that involves a change in production status (e.g. a research and development use case that will be promoted to a production environment);
 - ii. any change to the software release life cycle;
 - b. if the target audience for the AI use case changes (e.g. from internal to external users);
 - a. significant human-ai configuration changes (e.g. major changes in the content provided to users that may significantly change behavior);
 - c. if the use case's output type changes (e.g. from text to imagery);
 - d. if the solution architecture undergoes significant modification, including new system connections or process models;
 - e. major or minor update changes to underlying models as per [Semantic Versioning](#) standards; and

- f. Any other modifications that meet the definition of 'significant modification' put forth by the National Institute of Standards and Technology's Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy.
17. Underserved Communities: Has the meaning established in Section 10(b) of Executive Order 14091.
18. Use Case Register: a registry of all non-excluded AI use cases within GSA.

5. Appendix A: Presumed Rights-Impacting and Safety Impacting Use Cases

5.1 Rights-Impacting Use Case Examples

The following examples will be adjudicated as rights-impacting if used to control or meaningfully influence the outcomes of any of the following non-exhaustive list of activities or decisions:

- a. Blocking, removing, hiding, or limiting the reach of protected speech;
- b. In law enforcement contexts, producing risk assessments about individuals; predicting criminal recidivism; predicting criminal offenders; identifying criminal suspects or predicting perpetrators' identities; predicting victims of crime; forecasting crime; detecting gunshots; tracking personal vehicles over time in public spaces, including license plate readers; conducting biometric identification (e.g. iris, facial, fingerprint, or gait matching); sketching faces; reconstructing faces based on genetic information; monitoring social media; monitoring prisons; forensically analyzing criminal evidence; conducting forensic genetics; conducting cyber intrusions in the course of an investigation; conducting physical location-monitoring or tracking of individuals; or making determinations related to sentencing, parole, supervised release, probation, bail, pretrial release, or pretrial detention;
- c. Deciding or providing risk assessments related to immigration, asylum, or detention status; providing immigration-related risk assessments about individuals who intend to travel to, or have already entered, the U.S. or its territories; determining individuals' border access or access to Federal immigration related services through biometrics or through monitoring social media and other online activity; monitoring individuals' physical location for

- immigration and detention-related purposes; or forecasting the migration activity of individuals;
- d. Conducting biometric identification for one-to-many identification in publicly accessible spaces;
 - e. Detecting or measuring emotions, thought, impairment, or deception in humans;
 - f. Replicating a person's likeness or voice without express consent;
 - g. In education contexts, detecting student cheating or plagiarism; influencing admissions processes; monitoring students online or in virtual-reality; projecting student progress or outcomes; recommending disciplinary interventions; determining access to educational resources or programs; determining eligibility for student aid or Federal education; or facilitating surveillance (whether online or in-person);
 - h. Screening tenants; monitoring tenants in the context of public housing; providing valuations for homes; underwriting mortgages; or determining access to or terms of home insurance;
 - i. Determining the terms or conditions of employment, including pre-employment screening, reasonable accommodation, pay or promotion, performance management, hiring or termination, or recommending disciplinary action; performing time-on-task tracking; or conducting workplace surveillance or automated personnel management;
 - j. Carrying out the medically relevant functions of medical devices; providing medical diagnoses; determining medical treatments; providing medical or insurance health-risk assessments; providing drug-addiction risk assessments or determining access to medication; conducting risk assessments for suicide or other violence; detecting or preventing mental-health issues; flagging patients for interventions; allocating care in the context of public insurance; or controlling health-insurance costs and underwriting;
 - k. Allocating loans; determining financial-system access; credit scoring; determining who is subject to a financial audit; making insurance determinations and risk assessments; determining interest rates; or determining financial penalties (e.g. garnishing wages or withholding tax returns);
 - l. Making decisions regarding access to, eligibility for, or revocation of critical government resources or services; allowing or denying access—through

biometrics or other means (e.g. signature matching)—to IT systems for accessing services for benefits; detecting fraudulent use or attempted use of government services; assigning penalties in the context of government benefits;

- m. Translating between languages for the purpose of official communication to an individual where the responses are legally binding; providing live language interpretation or translation, without a competent interpreter or translator present, for an interaction that directly informs an agency decision or action; or
- n. Providing recommendations, decisions, or risk assessments about adoption matching, child protective actions, recommending child custody, whether a parent or guardian is suitable to gain or retain custody of a child, or protective actions for senior citizens or disabled persons.

5.2 Safety-Impacting Use Case Examples

The following examples will be adjudicated as safety-impacting if used to control or meaningfully influence the outcomes of any of the following non-exhaustive list of activities or decisions:

- a. Controlling the safety-critical functions within dams, emergency services, electrical grids, the generation or movement of energy, fire safety systems, food safety mechanisms, traffic control systems and other systems controlling physical transit, water and wastewater systems, or nuclear reactors, materials, and waste;
- b. Maintaining the integrity of elections and voting infrastructure;
- c. Controlling the physical movements of robots or robotic appendages within a workplace, school, housing, transportation, medical, or law enforcement setting;
- d. Applying kinetic force; delivering biological or chemical agents; or delivering potentially damaging electromagnetic impulses;
- e. Autonomously or semi-autonomously moving vehicles, whether on land, underground, at sea, in the air, or in space;
- f. Controlling the transport, safety, design, or development of hazardous chemicals or biological agents;
- g. Controlling industrial emissions and environmental impacts;
- h. Transporting or managing of industrial waste or other controlled pollutants;

- i. Designing, constructing, or testing of industrial equipment, systems, or structures that, if they failed, would pose a significant risk to safety;
- j. Carrying out the medically relevant functions of medical devices; providing medical diagnoses; determining medical treatments; providing medical or insurance health-risk assessments; providing drug-addiction risk assessments or determining access to medication; conducting risk assessments for suicide or other violence; detecting or preventing mental-health issues; flagging patients for interventions; allocating care in the context of public insurance; or controlling health-insurance costs and underwriting;
- k. Detecting the presence of dangerous weapons or a violent act;
- l. Choosing to summon first responders to an emergency;
- m. Controlling access to or security of government facilities; or
- n. Determining or carrying out enforcement actions pursuant to sanctions, trade restrictions, or other controls on exports, investments, or shipping.

6. Appendix B: AI Impact Statement Guidance

AI impact statements are necessary for any Safety-Impacting or Rights-Impacting use cases. A template for an impact statement may be found here:

[AI Impact Statement - Template](#). In AI impact statements, all system owners must document the following:

- a. The intended purpose for the AI and its expected benefit, supported by specific metrics or qualitative analysis. Metrics should be quantifiable measures of positive outcomes for the agency's mission – for example, to reduce costs, increase adoption, reduce wait time for customers, reduce risk to human life, or to meet compliance requirements – that can be measured using performance measurement or program evaluation methods after the AI is deployed to demonstrate the value of using AI. Where quantification is not feasible, qualitative analysis should demonstrate an expected positive outcome, such as for improvements to customer experience, and it should demonstrate that AI is better suited to accomplish the relevant task as compared to alternative strategies.
- b. The potential risks of using AI, as well as what, if any, additional mitigation measures, beyond these minimum practices, the agency will take to help reduce these risks. System owners should document the stakeholders who will be most

impacted by the use of the system and assess the possible failure modes of the AI and of the broader system, both in isolation and as a result of human users and other likely variables outside the scope of the system itself. System owners should be especially attentive to the potential risks to underserved communities. The expected benefits of the AI functionality should be considered against its potential risks, and if the benefits do not meaningfully outweigh the risks, system owners should not use the AI.

- c. The quality and appropriateness of the relevant data, or documentation on why those data are not available and what mitigations are in place. System owners must assess the quality of the data used in the AI's design, development, training, testing, and operation and its fitness to the AI's intended purpose. In conducting assessments, if the system owner cannot obtain such data after a reasonable effort to do so, it must obtain sufficient descriptive information from the vendor (e.g. AI or data provider) to satisfy the reporting requirements in this paragraph. At a minimum, system owners must document:
 - i. the data collection and preparation process, which must also include the provenance of any data used to train, fine-tune, or operate the AI;
 - ii. the quality and representativeness of the data for its intended purpose;
 - iii. how the data is relevant to the task being automated and may reasonably be expected to be useful for the AI's development, testing, and operation;
 - iv. whether the data contains sufficient breadth to address the range of real-world inputs the AI might encounter and how data gaps and shortcomings have been addressed either by the agency or vendor; and
 - v. if the data is maintained by the Federal Government, whether that data is publicly disclosable as an open government data asset, in accordance with applicable law and policy.

7. Appendix C: Additional Documents

- a. [Use Case Request Form](#)
- b. [AI Model Request Form](#)
- c. [AI Governance Board Charter](#)
- d. [AI Safety Team Charter](#)
- e. [Impact Statement Template](#)
- f. [Experimental Design Statement Template](#)