



U.S. General Services Administration (GSA)

GSA Order: Accelerating Responsible Use of Artificial Intelligence at GSA

CIO 2185.1C

GSA-IT

caio@gsa.gov

Purpose:

This directive updates the governing policies for the responsible, efficient, and accelerated adoption of artificial intelligence (AI) technologies and platforms within the General Services Administration (GSA). It aligns with the principles outlined in Office of Management and Budget (OMB) Memorandums M-25-21, M-25-22, and M-26-04 by promoting innovation, enhancing public trust, and ensuring mission-enabling use of AI. The directive establishes standards for the assessment, procurement, usage, monitoring, and governance of AI systems, emphasizing risk management, transparency, and lifecycle accountability. It supports the development of agency-wide AI strategies, fosters cross-functional collaboration, and encourages the use of trustworthy, interoperable, and American-made AI solutions. All activities under this directive must comply with existing security, privacy, and ethics regulations and applicable laws.

Background:

The AI in Government Act of 2020 (Public Law 116-260); AI Training Act of 2023 (Public Law 117-207); OMB memoranda M-25-21, M-25-22, and M-26-04; and OMB Circular No. A-119 direct all federal agencies to:

1. Accelerate the responsible adoption of AI by emphasizing innovation, governance, and public trust, in accordance with OMB M-25-21. Agencies must reduce bureaucratic barriers and promote mission-enabling AI that benefits the American public while safeguarding civil rights, civil liberties, and privacy.
2. Empower Chief AI Officers (CAIOs) and other AI leaders to drive strategic planning, workforce development, and cross-agency collaboration. Agencies must establish AI governance boards and publish public AI strategies that identify barriers and outline plans for scaling responsible AI use.
3. Ensure compliance with applicable federal laws and policies in the development, deployment, and use of AI and automated systems. This includes adherence to

privacy, safety, and nondiscrimination standards and alignment with voluntary consensus standards as outlined in OMB Circular A-119.

4. Establish and maintain processes to measure, monitor, evaluate, and report on AI use cases and their performance, especially for high-impact AI. Agencies must implement risk management practices and be prepared to terminate non-compliant systems.
5. Conduct regular AI risk assessments, particularly for high-impact systems, and integrate findings into governance and acquisition decisions. Agencies must also contribute to interagency repositories of best practices and tools.
6. Prioritize AI applications that advance the agency's mission, improve service delivery, and promote innovation. Agencies should invest in AI-enabling infrastructure, including data governance, workforce training, and reusable tools and models.
7. Ensure sufficient infrastructure and capacity for AI-ready data, including robust data curation, labeling, and stewardship practices. Agencies must manage data as a strategic asset to support trustworthy AI.
8. Assess and plan for AI workforce needs by identifying required competencies, offering training (as mandated by the AI Training Act), and aligning hiring and reskilling strategies with evolving AI demands.
9. Support interagency coordination and standards-setting initiatives, and encourage the adoption of voluntary consensus standards for AI, as directed by OMB Circular A-119 and M-25-21.
10. Drive efficient and responsible AI acquisition by fostering a competitive U.S. AI marketplace, managing performance and risk throughout the acquisition lifecycle, and ensuring cross-functional engagement in procurement decisions.

Applicability:

This order applies to:

1. All individuals, including GSA employees and contractors, who access, manage, share, or use data, including those involved in system-to-system data exchanges, and those participating in AI-related activities, training, or governance as defined under the AI Training Act of 2023.
2. All IT systems owned or operated by or on behalf of any GSA Service and Staff Office that process, store, or transmit federal data, especially where AI

capabilities are integrated or acquired, in accordance with OMB Memoranda M-25-21, M-25-22, and M-26-04.

3. All federal data contained in or processed by GSA IT systems, including data used to train, evaluate, or operate AI systems, subject to applicable privacy, security, and intellectual property (IP) safeguards outlined in M-25-22.
4. The Office of Inspector General (OIG) to the extent that participation aligns with the OIG's independent authority under the Inspector General Act of 1978 (5 U.S.C. App. 3) and does not conflict with OIG policies or its mission.
5. The Civilian Board of Contract Appeals (CBCA) only to the extent that participation is consistent with the CBCA's requisite independence under the Contract Disputes Act (41 U.S.C. §§ 7101–7109) and its legislative history.
6. All AI systems or services acquired by or on behalf of GSA, excluding common commercial products with embedded AI functionality not primarily used for AI purposes, as defined in OMB Memorandum M-25-22.

Standards and conformity assessments used in AI-related activities must align with OMB Circular A-119, which encourages the use of voluntary consensus standards and minimizes reliance on government-unique standards.

Cancellation:

This directive supersedes CIO 2185.1B, Use of Artificial Intelligence at GSA

Summary of Changes:

This revision makes changes to align with existing Executive Orders and OMB memoranda as referenced in the Background section.

Roles and Responsibilities:

1. **Chief AI Officer (CAIO):** In addition to the responsibilities defined in applicable Executive Orders and OMB memoranda, the CAIO must:
 - a. Maintain awareness of and promote the responsible and innovative use of AI technologies across GSA, including understanding system design, functionality, and intended use cases. Ensure AI adoption aligns with agency mission and public trust goals.

- b. Establish, maintain, and chair internal AI governance bodies, including the AI Governance Board and AI Oversight Committee, to ensure compliance with OMB M-25-21 and M-25-22.
 - c. Establish and update processes to measure, monitor, and evaluate the performance, accessibility, cost-effectiveness, and outcomes of AI applications, including risk management practices for high-impact systems.
 - d. Develop and oversee implementation of agency AI compliance plans, ensuring alignment with Executive Orders, OMB guidance, and applicable laws. Coordinate reporting to OMB, including AI use case inventories and acquisition strategies.
 - e. Ensure AI acquisitions comply with M-25-22 and M-26-04, including vendor sourcing, data portability, interoperability, privacy safeguards, truth-seeking and ideological neutrality. Collaborate with acquisition officials to embed appropriate contract terms and standards.
 - f. Collaborate with the Chief Human Capital Officer and Chief Learning Officer to identify skill gaps and implement training programs for federal managers and supervisors, as required by the AI Training Act of 2023.
 - g. Identify and convene federal and non-federal individuals, agencies, organizations, and entities with AI expertise to provide input relevant to GSA's mission functions, ensuring diverse perspectives and technical rigor.
 - h. In coordination with relevant officials, issue waivers for individual AI applications when appropriate. Establish and maintain criteria for categories of AI applications that do not require disposition through the AI Governance Board or AI Oversight Committee.
 - i. Promote the use of voluntary consensus standards in AI development and acquisition, consistent with OMB Circular A-119, and participate in relevant standards bodies where appropriate.
2. **AI Governance Board:** The Evidence-Based Data Governance Executive (EDGE) Board serves as GSA's principal AI governance body. In alignment with federal mandates and best practices, the EDGE Board shall:
- a. Define and periodically update the agency's vision, goals, and priorities for AI development and deployment, ensuring alignment with GSA's mission, values, and public service objectives.

- b. Support the CAIO in implementing governance frameworks, including risk management protocols, ethical guidelines, and compliance mechanisms for AI systems.
 - c. Ensure that AI systems used or acquired by GSA uphold principles of fairness, transparency, accountability, and privacy, consistent with OMB guidance and the AI Risk Management Framework.
 - d. Oversee the integration of AI-related competencies into workforce development plans, in coordination with the CAIO and Human Capital leadership, as required by the AI Training Act of 2023.
 - e. Convene relevant stakeholders across GSA—including legal, acquisition, IT, privacy, and program offices—to ensure coordinated oversight of AI initiatives and procurement.
 - f. Review agency-wide metrics and assessments of AI system performance, including risk mitigation strategies for high-impact use cases. Recommend corrective actions or system decommissioning when necessary.
 - g. Promote the use of voluntary consensus standards and conformity assessment practices in AI development and acquisition, consistent with OMB Circular A-119.
 - h. Validate the agency’s AI use case inventory, compliance plans, and reporting submissions to OMB, ensuring completeness, accuracy, and alignment with federal requirements.
 - i. Recommend external experts or advisory bodies to support GSA’s AI governance, particularly in areas of technical complexity, ethics, and emerging risks.
3. **AI Oversight Committee:** The AI Oversight Committee is a technical working group reporting to the CAIO in their role as co-chair of the EDGE Board. In alignment with federal mandates, the AI Oversight Committee shall:
- a. Implement the risk posture defined by the EDGE Board by developing and applying a standardized risk rubric for AI use cases. This includes managing intake, adjudication, and documentation of AI use cases across GSA.
 - b. Independently adjudicate Capability Assessment, GSA Pre-Approved Application, Research and Development, and Production or Production-Intent use cases.

- i. Production or Production-Intent use cases also require CAIO review and approval.
 - ii. High-impact use cases must be escalated to the EDGE Board for final adjudication.
- c. Deliver dispositions for Production or Production-Intent use cases, including assessments of feasibility, ethical implications, and compliance with legal and policy standards.
- d. Enforce all GSA-authorized security, privacy, and audit policies to protect Controlled Unclassified Information (CUI) and ensure AI systems operate within acceptable levels of residual risk. This includes:
 - 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)
 - vi. Federal Risk and Authorization Management Program (FedRAMP) authorizations
- e. Assess AI systems for performance, scalability, bias, transparency, and compatibility with existing infrastructure. Evaluate ethical implications, including fairness, accountability, and protection of individual rights.
- f. Include members with expertise in development, architecture, data science, user experience, privacy, security, and mission delivery. This ensures multi-disciplinary perspectives in AI governance.
- g. Maintain comprehensive records of adjudication decisions, rationale, and associated risks. Generate regular reports for internal review and external compliance reporting.
- h. Collaborate with system owners and executive sponsors to ensure ongoing monitoring of AI systems, especially those deemed high-impact. Recommend corrective actions or decommissioning if systems fail to meet standards.

- i. Engage with external stakeholders, including academic institutions, interagency councils, and civil society organizations, to benchmark practices and incorporate emerging standards.
4. **System Owner:** System owners shall:
- a. Register AI use cases with the AI Oversight Committee, including initial deployment and any significant modifications or decommissioning events. This supports centralized tracking and risk tiering as required by M-25-21.
 - b. Ensure compliance with this directive and all applicable OMB guidance, including privacy, security, and performance monitoring requirements for AI systems.
 - c. Support risk management by providing documentation and updates necessary for pre-deployment testing, ongoing monitoring, and human oversight of high-impact AI systems.
5. **Executive Sponsor:** Executive sponsors shall:
- a. Sponsor AI use cases in Capability Assessment, Research and Development, and Production or Production-Intent phases, ensuring alignment with strategic objectives and risk posture defined by the EDGE Board.
 - b. Champion responsible AI acquisition by coordinating with acquisition officials to ensure AI systems meet interoperability, privacy, and performance standards outlined in M-25-22.
 - c. Facilitate cross-functional engagement across legal, privacy, IT, and program offices to ensure AI initiatives are well-resourced and compliant with lifecycle governance requirements.
6. **Authorized Users of IT Resources:**
- a. General practitioners shall:
 - i. Protect federal nonpublic information and report any potential IT security incidents, consistent with GSA's IT General Rules of Behavior.
 - ii. Report unregistered or high-impact AI use cases to the AI Oversight Committee if they believe the system owner has not done so, supporting transparency and risk mitigation as required by M-25-21.

Table of Contents

- 1. INTRODUCTION..... 11**
 - 1.1. Objectives..... 11
 - 1.2. Scope..... 11
 - 1.3. Principles..... 11
- 2. POLICY..... 12**
 - 2.1. General AI usage..... 12
 - 2.2. New or Proposed AI Use Cases..... 14
 - 2.3. Existing AI Use Cases..... 15
 - 2.4. AI Code and Models..... 16
 - 2.5. Data Assets and Sources..... 17
 - 2.5.1. Internal Data Assets..... 18
 - 2.5.2. External Data Assets..... 18
 - 2.5.3. AI-Generated Data Products..... 19
 - 2.5.4. Data Dissemination Requirements..... 20
 - 2.6. Efficient Procurement of AI..... 21
 - 2.6.1. Pre-solicitation..... 21
 - 2.6.2. Procuring AI..... 22
 - 2.6.3. Procurement Policy Updates..... 23
 - 2.7. Tool or Product AI Enhancements..... 23
 - 2.8. Publication Requirements..... 24
 - 2.9. High-Impact AI..... 25
 - 2.9.1. Minimum Requirements for High-Impact AI..... 25
 - 2.9.2. Additional Requirements for Higher-Impact AI..... 26
 - 2.9.3. Excepted scenarios for Higher-Impact AI use cases..... 26
 - 2.9.4. Use-Case Waivers..... 27
 - 2.10. Organizational Risk Tolerance and Use Case Risk Rubric..... 27
- 3. DEFINITIONS..... 27**
- APPENDIX A: Presumed High-Impact Use Cases..... 33**
- APPENDIX B: AI Impact Statement Guidance..... 36**
- APPENDIX C: Additional Documents..... 38**

1. INTRODUCTION

As AI technologies evolve and expand within the federal government, their use must be managed to maximize effectiveness while minimizing potential harm and mitigating risks. AI has the potential to augment or improve mission delivery, service offerings, and productivity across all GSA equities. With appropriate oversight, controls, and human intervention protocols in place, AI can also promote innovation, productivity, and service delivery at scale. To capitalize on the potential benefits of these emerging technologies, policy controls must be established for the responsible, efficient, and accelerated development and use of AI.

This directive outlines the controls for AI usage within GSA, the governance and oversight infrastructure for responsible AI use, the processes for GSA employees to develop AI use cases, and the disclosure requirements for all AI implementations.

1.1. Objectives

The objectives of this order are to:

1. Define the AI governance model and procedures necessary to promote adoption of AI technologies while managing associated risks for GSA business activities, fostering a competitive American AI marketplace, and safeguarding taxpayer dollars;
2. Enable use of AI that improves service delivery, efficiency, and public trust in government;
3. Establish the roles, responsibilities, and reporting structures of the requisite oversight and governing groups, promoting cross-functional engagement in AI acquisition;
4. Outline the requirements of all AI systems, with noted focus on high-impact AI systems; and
5. Define core AI terms and concepts.

1.2. Scope

This order provides guidance for GSA program operations that have direct or indirect responsibility for, or control over, any action, activity, or program that relates to AI systems, including 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, transparency, safety, and collaboration. AI systems must be developed and deployed in a manner that prioritizes the public good while also considering 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. Capability Assessment: Explore and assess AI tools using only non-sensitive, public information and publicly available systems;
2. GSA Pre-approved Application: Select from a GSA pre-approved list of AI applications for low-stakes, low-impact scenarios with limited potential for harm or significant consequences;
3. Research and Development: Develop a capability using internal systems, processes, and data, without the immediate intent to promote the research output to a production environment or workflow (AI pilot projects fall into this category); and
4. Production or Production-Intent: Develop an AI capability with the immediate intention of using it to support business operations.

2.1. General AI usage

For all AI use cases, individuals acting on behalf of GSA must register proposed uses via the AI Use Case Request Form (see Appendix C). The AI Oversight Committee will assess these requests, identify their risk profiles, and adjudicate cases classified as Research and Development, or Production or Production-Intent. The AI Oversight Committee may request additional guidance from the CAIO and EDGE Board as necessary.

1. Publicly available, approved third-party AI endpoints and tools not currently sanctioned for use will be blocked from the GSA network and GFE devices unless approved per the following guidelines:
 - a. Access will be made available upon completion of GSA's AI Use Case Request Form, detailing intended usage and acknowledgment of the requirements of this directive and GSA's IT General Rules of Behavior.

- b. Beyond the tools available agency-wide, only endpoints and tools approved by the CAIO, the Chief Information Security Officer (CISO), and the EDGE Board will be made available. For chat bot use cases, users are required to use USAi, which is available to all GSA employees and contractors with a gsa.gov email address.
 - c. No output from publicly available products or tools may be introduced as a GSA production work product without approval from the EDGE Board.
 - d. 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 CUI, shall not be used as inputs (e.g., prompts or training data) to AI systems without approval from the AI Oversight Committee or EDGE Board.
2. Any work product output materially modified by, or solely produced by, generative AI (GenAI) 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:
 - a. Data;
 - b. Code;
 - c. Text (e.g., temporary and permanent records);
 - d. Applications (e.g., chatbots, recommendation engines);
 - e. Audio;
 - f. Imagery; and
 - g. Video.
3. All production systems using AI capabilities that provide direct interface with the public may include:
 - a. Notice and explanation of its services written in [plain language](#); and
 - b. Human alternatives or fallback options where practicable.
4. All AI software must have a valid ATO prior to use for Research and Development, and Production or Production-Intent use cases.

5. Any output from large language models 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 the code or content conforms to Section 508 of the Rehabilitation Act of 1973 and to the Section 508 technical standards for information and communication technology accessibility.

2.2. New or Proposed AI Use Cases

All AI use case requests must be submitted to the AI Oversight Committee via the AI Use Case Request Form (see Appendix C). If a model that is not currently authorized is being requested, the use case applicant must also submit an AI Model Request Form (see Appendix C). Research and Development use case requests shall also submit an Experimental Design Statement.

For all use case types, applicants must:

1. Provide basic information to identify each AI use case, including name, responsible agency, bureau/component, and email address;
2. Indicate whether the use case should be withheld from public reporting and its current stage of development (pre-deployment, pilot, deployed, or retired);
3. Classify the case as either (a) high-impact, (b) presumed high-impact, but determined not high impact or (c) not high-impact;
4. Describe the purpose and expected outcomes of the AI use case;
5. Specify the use case's topic area (e.g., Administrative Functions, Human Resources, Emergency Management, Energy & the Environment);
6. Specify the AI classification (e.g., GenAI, Agentic AI, Classical/Predictive Machine Learning);
7. Explain the problem it intends to solve and its expected benefits; and
8. Describe the AI system's outputs.

For Research and Development, and Production or Production-Intent use cases, applicants must:

1. Provide details about the AI use case's operational aspects, such as its operational or pilot start date, whether the system was purchased from a vendor, developed in-house, or both (with a required vendor name if applicable);

2. Indicate the presence of an ATO and the associated system name;
3. Provide information about the data and code used in the AI use case, including a description of data used for training, fine-tuning, and evaluation;
4. Indicate whether the data is publicly disclosed, if PII is involved, if a PIA is available, and which demographic variables are explicitly used as model features; and
5. Indicate whether custom-developed code is included and, if open source, include a link to the source code.

For use cases predetermined to be high-impact, applicants must:

1. Provide information about pre-deployment testing, completion of an AI impact assessment (including potential impacts and how they were identified), independent review, ongoing monitoring processes, human training for operators, the presence of fail-safes, and established appeal processes; and
2. Describe steps taken to consult and incorporate feedback from end users and the public.

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

2.3. Existing AI Use Cases

Every year, existing Research and Development, or Production or Production-Intent use cases must be re-registered with the AI Oversight Committee via the AI Use Case Request Form (see Appendix C). The same form shall be used to report retired use cases. The specific requirements are as follows:

1. All existing Research and Development, or Production or Production-Intent use cases shall be reported to the AI Oversight Committee on an annual basis.
2. Any use case that undergoes a significant modification must be re-submitted to the AI Oversight Committee for reassessment. Any modification that elevates the use case into a high-impact AI use case will be re-adjudicated by the AI Oversight Committee.
3. Any use case where there has been a cybersecurity or privacy incident must be re-submitted to the AI Oversight Committee for re-evaluation within 5 days of the reported incident.

4. AI systems that use Federal nonpublic information shall be conducted within approved secure enterprise systems, such as the Enterprise Data Solution (EDS).
5. 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:
 - a. CAIO;
 - b. EDGE Board;
 - c. The AI Oversight Committee;
 - d. CISO; or
 - e. 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.

1. All internally developed AI code, models, and model weights shall be:
 - a. Shared within GSA for internal reuse and learning; and
 - b. Open-sourced to public repositories, consistent with GSA's Open Source Software (OSS) Policy, unless an exception applies.
2. AI code and models that are no longer in active use may be archived. Maintenance is not required for archived assets unless they are reactivated for operational use or reference.
3. Use of open source or commercial-off-the-shelf (COTS) models shall:
 - a. Be reviewed and approved by the AI Oversight Committee;
 - b. Be treated as a system integration activity, subject to GSA's enterprise architecture standards;
 - c. Comply with applicable security, privacy, and interoperability requirements, including those outlined in OMB M-25-22 and Circular A-119.
4. Code and models may be exempt from public release if:

- a. Disclosure is restricted by law, regulation, or Executive Order;
- b. Sharing poses identifiable risks to national security, privacy, public safety, or the confidentiality of government information;
- c. Assets were developed solely for research and development purposes;
- d. Contractual obligations prohibit sharing; or
- e. Disclosure would compromise agency mission, operations, or system integrity.

2.5. Data Assets and Sources

To ensure transparency, accountability, and trust in GSA's use of AI, all system owners of Research and Development, Production, or Production-Intent AI systems must document and report the data used throughout the AI lifecycle. This includes data used in the design, development, training, testing, and operation of AI systems.

System owners must report:

1. The types and sources of data used within the AI system;
2. The intended purpose of each dataset;
3. The data owners or stewards;
4. The relevance of the data to the automated task or decision; and
5. The sensitivity level of the data, including any CUI

In alignment with OMB M-25-22, the AI Oversight Committee or CAIO may request additional documentation, including:

1. Data provenance and preparation processes;
2. Quality, completeness, and representativeness of the data;
3. Intended use of the data across the AI lifecycle; and
4. Evidence that the data adequately reflects real-world conditions, and mitigation strategies for any gaps or biases.

All data used in AI systems must:

1. Be registered and published in the Enterprise Data Catalog in EDS, in accordance with GSA's data governance policies;

2. Comply with the Internal Data Sharing Directive, including its data categorization and sensitivity framework; and
3. Adhere to OMB Circular A-119 by prioritizing the use of open, interoperable, and voluntary consensus standards for data formats and metadata schemas.

Exceptions to data disclosure and cataloging may apply when:

1. Disclosure is restricted by law, regulation, or Executive Order;
2. Sharing poses identifiable risks to national security, privacy, or public safety;
3. Data is used exclusively for exploratory research and development;
4. Contractual obligations prohibit disclosure; or
5. Disclosure would compromise agency mission, operations, or system integrity.

2.5.1. Internal Data Assets

To ensure responsible use of internal data in AI systems and uphold federal standards for privacy, security, and trustworthiness, GSA mandates the following requirements for internal data assets used in AI development and deployment:

1. Internal data assets shall not be used as input to public-facing AI systems or services unless explicitly authorized. This restriction is in place to prevent unauthorized disclosure, misuse, or unintended exposure of government data.
2. No CUI may be used in any AI system without:
 - a. Clearance from the AI Oversight Committee;
 - b. Submission and approval of an AI Impact Statement, as required under OMB M-25-22 for high-impact AI systems; and
 - c. A valid ATO issued through the appropriate cybersecurity and privacy review channels.

2.5.2. External Data Assets

To ensure the integrity, traceability, and responsible use of externally sourced data in AI systems, GSA requires that all external data assets used in Research and Development, Production, or Production-Intent AI systems be documented and governed as follows:

1. The originator, collection methodology, and preparation process for all external data sources must be registered with the AI Oversight Committee. This documentation must be:
 - a. Submitted during initial system review;
 - b. Resubmitted annually to maintain approval for continued use; and
 - c. Updated immediately if any significant modifications are made to the AI use case or the data source.
2. All approved external data sources shall be:
 - a. Indexed and maintained in the Enterprise Data Catalog in EDS; and
 - b. Made accessible to relevant stakeholders for review, audit, and reuse, consistent with GSA's data governance policies.
3. External data must comply with:
 - a. Federal privacy, security, and ethical standards, including those outlined in OMB M-25-22 for trustworthy AI;
 - b. GSA's internal data classification and sharing directives; and
 - c. OMB Circular A-119, which encourages the use of voluntary consensus standards for data formats, metadata, and interoperability.

2.5.3. AI-Generated Data Products

To ensure transparency, traceability, and responsible use of AI-generated outputs, all data products created or modified by AI systems must be clearly labeled and governed in accordance with federal standards.

1. All AI-generated data outputs—including original datasets, augmented records, and modified fields—must:
 - a. Be labeled as AI-generated in their metadata;
 - b. Include the name and version of the AI system responsible for the generation or modification; and
 - c. Be indexed and cataloged in the Enterprise Data Catalog in EDS for internal discovery and audit purposes.
2. For datasets that have undergone AI-driven augmentation (e.g., imputation, field creation, enrichment), metadata must specify:

- a. Which records or fields were modified or created;
 - b. The nature of the augmentation; and
 - c. The AI system and version used in the process.
3. All AI-generated data must comply with:
 - a. GSA's existing data governance, privacy, and security policies;
 - b. Federal standards for trustworthy AI as outlined in OMB M-25-22, including documentation of provenance and system accountability; and
 - c. OMB Circular A-119 guidance on metadata standards and interoperability.
4. Certain AI-generated metadata elements may be exempt from explicit labeling, including:
 - a. Field titles, descriptions, and domain associations used for metadata enrichment;
 - b. Classification or tagging labels used for search optimization or discovery;
 - c. Domain associations used for ontology management; and
 - d. Human-authored content (e.g., emails, chats) that includes auto-completed or suggested text, provided the AI contribution does not materially alter authorship or intent.

2.5.4. Data Dissemination Requirements

To promote transparency, accountability, and public access to government data, all datasets used in the development, training, testing, or evaluation of AI models or applications must be treated as public data assets unless restricted by law or policy.

1. Data used in AI systems shall be evaluated and, where appropriate, designated as Open Government Data Assets under the Open, Public, Electronic, and Necessary (OPEN) Government Data Act. These assets must:
 - a. Be made publicly available via Data.gov;
 - b. Include appropriate metadata and documentation to ensure usability and discoverability; and
 - c. Comply with the Foundations for Evidence-Based Policymaking Act of 2018 and OMB Memorandum M-25-05, which outlines open data access and management requirements.

2. All data products identified for dissemination must:
 - a. Undergo a privacy and risk assessment to ensure compliance with applicable laws and policies, including those governing PII or other CUI; and
 - b. Be reviewed by the AI Oversight Committee and relevant data governance bodies to ensure that dissemination does not compromise national security, individual privacy, or agency operations.
3. Disseminated data must adhere to:
 - a. GSA's internal data standards and classification frameworks;
 - b. OMB Circular A-119 guidance on the use of voluntary consensus standards for data formats, metadata, and interoperability; and
 - c. Any applicable federal or agency-specific data quality, accessibility, and documentation requirements.

2.6. Efficient Procurement of AI

2.6.1. Pre-solicitation

To ensure that AI technologies are responsibly considered and integrated into GSA-funded procurements, acquisition teams must incorporate AI-related planning and oversight during the early stages of market research and acquisition strategy development.

1. During market research, acquisition teams must assess whether AI capabilities may be proposed as part of a vendor's solution. If AI is identified as a potential component of the procurement, teams must coordinate the acquisition plan and requirement with the CAIO and AI Oversight Committee in accordance with Chapter 4: Technology of the GSA Acquisition Handbook. The CAIO and AI Oversight Committee will:
 - a. Evaluate risks, benefits, and alignment with GSA's AI governance framework;
 - b. Ensure alignment to OMB M-25-22 and M-26-04, including documentation, transparency, and human oversight;
 - c. Assist with considering the implications of the AI Training Act of 2023, ensuring that acquisition professionals are equipped to evaluate AI proposals;

- d. Assist with aligning the requirement, as applicable, with the AI in Government Act of 2020, which encourages responsible AI adoption and interagency collaboration; and
- e. Ensure that any AI systems acquired adhere to voluntary consensus standards for safety, interoperability, and performance reference in OMB Circular A-119.

2.6.2. Procuring AI

In accordance with Chapter 4: Technology of the GSA Acquisition Handbook, prior to the release of a solicitation for AI for use at GSA, the acquisition team must ensure the requirements document (Performance Work Statement, Statement of Objective, or Statement of Work) has been coordinated and approved by the CAIO. If the AI system will create, store, or otherwise use PII, the Senior Agency Official for Privacy needs to have early and ongoing involvement in the acquisition to manage privacy risks and ensure compliance with law and policy related to privacy.

Contracting activities must maximize buying commercial products or commercial services, including AI, rather than requiring Government-unique solutions, in accordance with FAR 1.102(a)(3) and Executive Order 14271, *Ensuring Commercial, Cost-Effective Solutions in Federal Contracts* and must procure AI products and services from existing governmentwide contracts (i.e., Federal Supply Schedule) to the maximum extent practicable.

Submit all requirement documents to the AI Oversight Committee. Solicitations cannot be released without coordination with the CAIO. Any solicitation involving a planned or likely high-impact AI system use case must disclose this fact.

AI system or service contracts should clearly define data ownership and IP rights, aiming for contract standardization. These contracts must:

1. Scope licensing and IP rights to prevent vendor lock-in, based on AI's intended use;
2. Guarantee the acquiring agency retains access to necessary AI system components for operation and monitoring;
3. Guide vendor handling of agency data, ensuring collection and retention only when contractually necessary;
4. Prohibit the permanent use of non-public agency data (input/output) for training public or commercial AI without explicit agency consent, as per applicable law; and

5. Prioritize documentation for transparency, explainability, and effective performance tracking of procured AI.

2.6.3. Procurement Policy Updates

GSA's procurement policies for AI technologies will be iterated to reflect evolving federal standards, updates to the Federal Acquisition Regulation and GSAR/M requirements, agency directives, and best practices for responsible AI acquisition.

1. All updates will maintain compliance with updates to the Federal Acquisition Regulation (FAR) and the GSAR/M, ensuring that AI-related procurements are conducted with transparency, fairness, and accountability.
2. Updates to procurement policy will be coordinated with the CAIO and AI Oversight Committee to ensure that all AI-related acquisitions are subject to appropriate governance, risk assessment, and ethical review. Updates to GSA AI directives will be coordinated with the GSA Office of Acquisition Policy (OGP) to ensure compliance with the FAR updates and GSAR/M.
3. GSA collects and applies lessons learned for AI acquisitions in alignment with OMB Memorandums M-25-21 and M-25-22 through a combination of structured procurement strategies, stakeholder engagement, and iterative feedback mechanisms.

2.7. Tool or Product AI Enhancements

To ensure responsible integration of AI capabilities into existing GSA systems and tools, any AI enhancements applied to products with an active ATO must undergo review and reauthorization in accordance with federal and agency-specific AI governance policies.

1. If an existing GSA tool or product receives an AI enhancement:
 - a. The system owner must submit the updated system to the Authorizing Official (AO) for assessment; and
 - b. A revised ATO must be issued prior to deploying the AI functionality within the system's operational boundary.
2. AI enhancements must be:
 - a. Reported as a procurement activity via the AI Use Case Request Form (see Appendix C); and

- b. Reviewed and dispositioned by the AI Oversight Committee to ensure compliance with GSA's AI directive and federal standards for trustworthy AI.
- 3. If the AI enhancement violates any provision of this directive or applicable federal guidance (e.g., OMB M-25-22), the system owner must:
 - a. Disable the enhancement or revert to a prior version of the tool that does not include the AI functionality; or
 - b. If reversion is not feasible, submit compensating process controls and policy documentation to the CAIO demonstrating that the enhancement is not used in prohibited or high-risk use cases.

2.8. Publication Requirements

To promote transparency, public trust, and compliance with federal mandates for open government and responsible AI use, GSA shall publicly disclose information about its AI systems in accordance with applicable laws and guidance.

- 1. All Research and Development, and Production and Production-Intent AI systems currently in use must be:
 - a. Included in GSA's AI Use Case Inventory;
 - b. Published on gsa.gov or another designated public-facing platform; and
 - c. Accompanied by the data elements required by OMB's integrated data collection process, or any superseding process designated by OMB, as outlined in OMB M-25-05.
- 2. AI systems may be exempt from public disclosure if:
 - a. Disclosure would conflict with applicable law, regulation, or government-wide policy; or
 - b. Public release would pose a demonstrable risk to national security, privacy, or operational integrity.
- 3. GSA shall publish aggregate statistics on its AI use cases, including:
 - a. The number of high-impact AI use cases currently in operation;
 - b. The compliance status of all AI systems with applicable directives and standards; and

- c. A list of approved waivers or exceptions currently in force.

2.9. High-Impact AI

2.9.1. Minimum Requirements for High-Impact AI

All AI use cases that the AI Oversight Committee determines are high-impact AI may be subject to the additional requirements in this section because of the potential risk they can pose, for example, harms to people or businesses. Recognizing both the risks and opportunities presented by potential higher-impact 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.

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

1. All high-impact use cases may be required to 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; and
 - c. Submit to an independent evaluation of the AI system from the CAIO or their designee.
2. All high-impact use cases may be required to follow these practices *while* employing the AI into any use case:
 - a. Conduct ongoing monitoring of the AI system and establish thresholds for periodic human review;
 - b. Mitigate emergent risks identified through routine testing, continuous monitoring protocols, or third-party findings;
 - c. Ensure all system practitioners have taken requisite AI training requirements;
 - d. Include human validation and intervention protocols to ensure all output decisions made by AI systems are regularly evaluated by system practitioners; and

- e. Provide public notice and plain language documentation regarding the high-impact use case through the public interface, in public disclosure statements, and in the AI use case inventory.

2.9.2. Additional Requirements for Higher-Impact AI

AI use cases deemed high-impact must follow these additional requirements *before* implementation:

1. Proactively identify and mitigate algorithmic discrimination or bias;
2. Assess and mitigate disparate impacts for protected classes;
3. Conduct direct user testing of system interactions; and
4. Solicit comments from the user community and conduct post-transaction customer feedback activities.

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

1. Conduct ongoing monitoring studies for AI-enabled discrimination;
2. Notify any negatively affected individuals;
3. Provide fallback and escalation options for AI processes or outcomes; and
4. Provide opt-out alternatives where practicable.

2.9.3. Excepted scenarios for Higher-Impact AI use cases

The following high-impact AI use cases do not need to follow the requirements set out in 2.9.1 and 2.9.2 above:

1. 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;
2. Evaluation of a particular AI application because the AI provider is the target or potential target of a regulatory enforcement action; and
3. Research and development purposes.

2.9.4. Use-Case Waivers

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 overall or would create an unacceptable impediment to critical agency operations.

1. Appeals for waivers may be submitted by system owners or delegates with written justifications to the CAIO.
2. All waivers must be centrally tracked and are subject to publication requirements outlined in the [Publication Requirements](#) section.
3. All waivers will be reassessed on an annual basis.

2.10. Organizational Risk Tolerance and Use Case Risk Rubric

The EDGE 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 Board](#) reporting requirements, under [GSA's Enterprise Risk Management Policy](#). This includes:

1. Establishing likelihood and impact ranking criteria and thresholds;
2. Defining the considered factors for the use case risk rubric; and
3. Establishing the risk management practices and processes required for AI systems.

The AI Oversight Committee is responsible for assessing use cases based on the guidance provided by the EDGE Board.

Each system owner is responsible for implementing the risk management processes defined by the EDGE Board.

3. DEFINITIONS

1. American AI Systems: The terms American AI Systems means AI systems developed and produced in the United States. (Office of Management and Budget Memorandum M-25-22).
2. 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.
3. AI System: The term "AI System" has the same meaning as defined in the Advancing American AI Act, section 7223(4) of Pub. L. 117-263.

4. AI Use Case: The application of AI 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 include using machine learning algorithms to optimize procurement processes, or leveraging natural language processing to improve search functionality on the agency's website.
5. Contestability: The ability to effectively challenge a decision made or augmented by AI.
6. Covered AI: AI that has been adjudicated to be high-impact.
7. Controlled Unclassified Information (CUI): Information the Government creates or possesses, or that an entity creates or possesses for or on behalf of the Government, that a law, regulation, or Government-wide policy requires or permits an agency to handle using safeguarding or dissemination controls.
8. Data Asset: a collection of data elements or data sets that may be grouped together.” (44 U.S.C. § 3502).
9. Federal Information: “The term “Federal information system” means an information system used or operated by an executive agency, by a contractor of an executive agency, or by another organization on behalf of an executive agency.” (40 USC § 11331(g)(1)).
10. Large Language Model (LLM): LLM means a generative AI model trained on vast, diverse datasets that enable the model to generate natural-language responses to user prompts. (Executive Order 14319 Section 2).
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; and
 - f. Other deliverables: External communications, training materials, marking collateral, and audit findings.
12. Research and Development: “Research and experimental development activities are 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.” (OMB Circular No. A-11).
13. Risks from the Use of AI: Risks related to efficacy, safety, 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.
- a. This includes such risks regardless of whether:
 - i. The AI merely informs the decision or action, partially automates it, or fully automates it;
 - ii. There is or is not human oversight for the decision or action;
 - iii. 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
 - iv. 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.
 - b. While the particular forms of these risks continue to evolve, at least the following factors can create, contribute to, or exacerbate these risks:
 - i. AI outputs that are inaccurate or misleading;
 - ii. AI outputs that are unreliable, ineffective, or not robust;
 - iii. AI outputs that are discriminatory or have a discriminatory effect;
 - iv. 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;

- v. AI being used for tasks to which it is poorly suited or being inappropriately repurposed in a context for which it was not intended;
- vi. 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
- vii. The adversarial evasion or manipulation of AI, such as an entity purposefully inducing AI to misclassify an input.

14. 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.

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 promoted to a production environment); or
 - ii. Any change to the software release lifecycle.
- b. Target Audience: If the target audience for the AI use case changes (e.g., from internal to external users).
- c. Human-AI Configuration: Significant human-AI configuration changes (e.g., major changes in the content provided to users that may significantly alter behavior).
- d. Output Type: If the use case's output type changes (e.g., from text to imagery).
- e. Solution Architecture: If the solution architecture undergoes significant modification, including new system connections or process models.
- f. Underlying Models: Major or minor update changes to underlying models as per Semantic Versioning standards.

- g. Other Modifications: Any other modifications that meet the definition of 'significant modification' as defined 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.
15. 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 or information system security managers. System owners represent the interests of the system throughout its lifecycle. Primary responsibility for managing risk should rest with the system owners.
 16. Use Case Register: A registry of all non-excluded AI use cases within GSA.

APPENDIX A: Presumed High-Impact Use Cases

This appendix provides a non-exhaustive list of AI use cases that are presumed to be high-impact. Final determinations regarding a high-impact AI use case will be based on the broader definition provided in this GSA Order.

The following examples will be adjudicated as high-impact if used to control or meaningfully influence the outcomes of any of the following activities or decisions:

1. Safety-Critical Functions: Controlling safety-critical functions within critical infrastructure systems or components, government facilities, emergency services (including calling 911), fire and life safety systems, food safety mechanisms, traffic control systems, water and wastewater systems, or nuclear reactors, materials, and waste.
2. Physical Movements: Controlling the physical movements of robots, robotic appendages, vehicles (on land, underground, at sea, in the air, or in space), or industrial equipment with the potential for significant human injury or death.
3. Kinetic or Non-Kinetic Measures: Applying kinetic force, delivering biological or chemical agents, delivering potentially damaging electromagnetic impulse, or other kinetic or non-kinetic measures for attack or active defense.
4. Hazardous Materials: Controlling the transport, safety, design, development, or use of hazardous chemicals or biological agents.
5. Equipment or Infrastructure Failure: Designing, constructing, or testing of industrial equipment, systems, or structures that, if they failed, would pose a significant risk to safety.
6. Healthcare Contexts: Carrying out the medically relevant functions of medical devices, providing patient 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.
7. Access or Security of Government Facilities: Controlling access to or security of government facilities.
8. Sanctions and Trade Restrictions: Determining or carrying out enforcement actions pursuant to sanctions, trade restrictions, or other controls on exports, investments, or shipping.

9. Protected Speech: Blocking, removing, hiding, or limiting the reach of protected speech.

10. Law Enforcement Contexts:

- a. Producing risk assessments about individuals;
- b. Predicting criminal recidivism;
- c. Predicting criminal offenders;
- d. Identifying criminal suspects or predicting perpetrators' identities;
- e. Forecasting crime;
- f. Tracking personal vehicles over time in public spaces (e.g., using license plate readers);
- g. Conducting biometric identification (e.g., iris, facial, fingerprint, or gait matching);
- h. Sketching faces;
- i. Reconstructing faces based on genetic information;
- j. Monitoring social media;
- k. Conducting digital forensic techniques;
- l. Conducting cyber-interventions in the course of an investigation;
- m. Conducting physical location-monitoring or tracking of individuals;
- n. Detecting the presence of dangerous weapons or violent activity; or
- o. Making determinations related to sentencing, parole, supervised release, probation, bail, pretrial release, or pretrial detention.

11. Immigration, Asylum, or Detention: Deciding or providing risk assessments related to immigration, asylum, or detention status, including for individuals who intend to travel to, or have already entered, the U.S. or its territories.

12. Biometric Identification in Publicly Accessible Spaces: Conducting one-to-many biometric identification in publicly accessible spaces.

13. Critical Federal Services and Benefits: Making decisions regarding access to, eligibility for, or revocation of critical federal services, processes, and benefits

(e.g., federal loans, public housing); determining continued eligibility for such services or benefits; allowing or denying access—through biometrics or other means (e.g., signature matching)—to IT systems for accessing services or benefits; detecting fraudulent use or attempted fraudulent use of government services; or assigning penalties in the context of government benefits.

14. Federal Employment: Determining the terms or conditions of federal employment, including pre-employment screening, reasonable accommodation, pay, promotion, performance management, hiring, termination, recommending disciplinary action, or reassignment of workers.
15. Language Translation: Translating between languages (foreign and audiovisual) for the purpose of official communication to an individual where the responses are legally binding or directly inform agency decisions or actions.

APPENDIX B: AI Impact Statement Guidance

AI Impact Statements are necessary for any high-impact 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:

1. 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.
2. 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.
3. 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:
 - a. The data collection and preparation process, which must also include the provenance of any data used to train, fine-tune, or operate the AI;
 - b. The quality and representativeness of the data for its intended purpose;

- c. 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;
- d. 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
- e. 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.

APPENDIX C: Additional Documents

1. [AI Model Request Form](#): If a model that is not currently authorized is being requested, the use case applicant must also submit an AI Model Request Form. This template is divided into four main sections: requestor information, model information, model specifications, and performance metrics.
2. [AI Oversight Committee Charter](#): This GSA charter establishes the AI Oversight Committee, which facilitates the effective integration and utilization of AI technologies across GSA, ensuring AI initiatives align with organizational goals, promote accessibility, enhance content quality, leverage research and data, provide training, optimize design/user experience, and are supported by robust technology and infrastructure.
3. [AI Use Case Request Form](#): Applicants must register every proposed AI use case with this form, no matter what type (i.e., Capability Assessment, GSA Pre-approved Application, Research and Development, or Production or Production-Intent). Only the Research and Development, and Production or Production-Intent requests will be assessed by the AI Oversight Committee for approval.
4. [EDGE Board Charter](#): This GSA charter establishes the EDGE Board, integrating the responsibilities of the former AI Governance Board. The EDGE Board, sponsored by GSA IT, oversees and coordinates the agency's AI and data governance activities.
5. [Experimental Design Statement Template](#): Required for Research and Development use case requests, this structured template includes sections covering all necessary aspects of a research proposal, including objectives, methodology, ethical considerations, timelines, and budget.
6. [Impact Statement Template](#): Required for use cases classified as high-impact, this template is used to document potential risks, including those to underserved communities, and proposed mitigation strategies. It asks for identification of impacted stakeholders, assessment of possible failure modes, and a detailed explanation of the quality, appropriateness, collection process, representativeness, relevance, and coverage of the data used in the AI's design, development, training, testing, and operation.