United States Department of the Interior

OFFICE OF THE SECRETARY
Washington, DC 20240

February 18, 2026

Memorandum

To: Heads of Bureaus and Offices

Through: Paul A. McInerny
Chief Information Officer
Office of the Chief Information Officer

PAUL MCINERNY Digitally signed by PAUL MCINERNY
Date: 2026.02.18 15:48:24 -05'00'

From: Jay McMaster
Chief Artificial Intelligence Officer
Office of the Chief Information Officer

PETER DOUCETTE Digitally signed by PETER DOUCETTE
Date: 2026.02.18 15:22:46 -06'00'

Subject: Policy on Use of Generative Artificial Intelligence (GenAI)

1. Purpose

The intent of this policy is to complement, and not supersede, guidance provided by OMB Memorandum M-25-21 [1], U.S. Department of the Interior (Department, DOI, Interior) Secretarial Order 3444 [2], the Department's Artificial Intelligence (AI) Strategy [3], and the DOI Compliance Plan [4]. The Policy establishes terms and guidelines for the responsible use of generative artificial intelligence (GenAI) technologies within the Department. Given the pace of GenAI innovation, a goal of this policy is to incorporate flexibility to accommodate clarifications and unanticipated changes over time, while ensuring responsible use.

2. Background

GenAI, defined in Appendix section A.1, is rapidly reshaping many aspects of society, including how federal agencies fulfill mission work. Large language models (LLMs), defined in Appendix section A.1, are the most widely used form of GenAI. While these technologies offer significant opportunities to improve efficiency, enhance decision-support, and streamline routine tasks, LLMs introduce new risks that must be managed responsibly. It is every employee's responsibility to understand how to use GenAI tools safely, ethically, and in a manner that upholds federal standards for security, privacy and quality.

3. Scope

This policy applies to all employees, contractors, and affiliates of the Department who use GenAI tools in the course of their official duties.

4. Secure and Trusted Use

GenAI tools may be applied to use cases involving both sensitive (private) and non-sensitive (public) data. Sensitive data requires trusted protection because unauthorized access could cause harm to individuals, organizations, or operations. Understanding the level of trust in the environment where GenAI tools are used is essential. A "trust boundary" acts like a security fence that manages access to data. A data leak occurs when data is shared or processed outside of an approved trust boundary, which risks unauthorized access. Enterprise-licensed GenAI tools, such as those provisioned through the Department's managed IT environments, are generally deployed within private trust boundaries that meet industry-grade security standards. By contrast, use of GenAI tools within public environments transmit data over public networks that do not provide the same level of trust assurance and risk data leakage. Subsections 4.1-4.3 address trust boundary scenarios.

4.1 Use of GenAI Tools with Sensitive Data

Only GenAI tools that operate within approved trust boundaries are authorized for use with sensitive data within the Department. A bureau or office executive leadership team with IT management responsibility should be consulted to clearly determine the trust boundary within which a GenAI tool is operated. Users of GenAI tools are responsible for determining data sensitivity based on document markings or by seeking appropriate guidance. Examples of sensitive data include, but are not limited to, personally identifiable information (PII), pre-decisional information, procurement-sensitive materials, or controlled unclassified information (CUI).

4.2 Use of GenAI Tools Over Public Networks

Use of GenAI tools over public networks is permitted provided they are 1) not used with sensitive data, 2) not prohibited by supplemental guidance, and 3) used in a manner consistent with the IT Rules of Behavior for Computer Network Users [5]. Examples of public network GenAI tools include, but are not limited to: personal free or subscription-based accounts with LLM providers; public web search tools provided within LLMs (see section 4.3), and software applications that use LLMs over public networks. Best practice is to prioritize use of GenAI tools that operate within agency-approved trust boundaries.

4.3 Data Handling in Work (Private) Versus Web (Public) Modes

GenAI tools enable users to augment pretrained LLM content with the addition of customized data sources. For example, "work" modes that integrate prompts with personal email and files must be authorized and configured for use with sensitive data. In contrast, "web" modes that transmit data from LLM prompts over public networks to search or access internet content are not authorized for use with sensitive data.

5. Human Oversight and Judgement (Human in the Loop)

Even the most secure operating environments cannot protect users from misusing GenAI output. Users need to understand the inherent risks when using GenAI and manage expectations accordingly. Responsible use includes understanding how output is generated (refer to Appendix section A.2), how to effectively prompt LLMs to reduce the potential for erroneous output (hallucinations), and how to apply appropriate judgement when interpreting GenAI output. Examples include:

5.1 Accuracy

LLMs generate output based on learned patterns from the data on which they are trained. Hallucinations—fabricated outputs—can occur when an LLM lacks sufficient patterns to make reliable estimates. Because an LLM is unaware of these gaps, it will provide its best estimate, which may appear confident but can be inaccurate. Best practices to detect or reduce hallucinations include asking the LLM for references it used, providing sufficient context in prompts, and validating responses against independent sources.

5.2 Variability

GenAI systems produce probabilistic outputs by design, meaning the same input may yield different output variations. This behavior is useful to increase creativity. While configuration settings can reduce variability in some tools, more consistent output does not guarantee its reliability. Output accuracy must be validated with human oversight.

5.3 Bias

LLM responses reflect the data they were trained on, which can include unexpected or subtle biases. Because these biases are hard to detect and remove, it is important to reduce their impact. Best practices include critically reviewing outputs, framing prompts carefully, and using comprehensive or iterative prompts supported by reliable sources.

6. Risk Management

Responsible risk management is a deliberate and unbiased approach to weighing risk proportionately. For example, natural tensions are inherent between goals of safety and acceleration of GenAI use. Differentiating among high-impact and all other use cases, as defined in OMB M-25-21 [1], allows managers to tailor risk *minimization* and risk *acceptance* decisions that are consistent with objectives to accelerate the responsible use of GenAI. Additional best practices include:

- Clearly identifying actions proposed and/or taken that constitute risk minimization or risk acceptance.
- Exploring reward mechanisms to incentivize responsible risk acceptance that accelerates progress.
- Further identifying and managing the unique risks posed by GenAI systems, such as hallucination, misuse, bias, and lack of traceability by consulting the National Institute of Standards and Technology (NIST) AI Risk Management Framework: Generative AI Profile [6].

7. Transparency and Public Trust

Public trust in GenAI systems depends on clear communication, accountability, and mechanisms for oversight. Responsible use includes disclosing GenAI use when required, establishing mechanisms for public feedback, and ensuring transparency in the use and governance of GenAI tools.

7.1 Citing Use of GenAI

Official content generated or significantly influenced by GenAI should include an appropriate level of citation, such as "This document was drafted with the assistance of [GenAI Tool X, month, year]." Official content includes, but is not limited to, any material that represents an organization's position, findings, or decisions that are shared for internal deliberation or external dissemination. GenAI users should be mindful that content publishers such as science journals may impose additional requirements for citing the use of GenAI tools.

7.2 Feedback Mechanisms

Channels must be available to communicate GenAI use to the public and allow feedback on GenAI-generated content and tool governance. These mechanisms ensure accountability and identify unintended consequences or misuse. Examples include public-facing websites that solicit feedback and periodic open reporting of activities.

7.3 FOIA Requests

GenAI prompts, generated content, and related decision-making processes in support of official duties may be subject to Freedom of Information Act (FOIA) requests. Best practices include maintaining records of GenAI use in support of official duties and ensuring that such records are accessible and reviewable by the public.

8. Training and Education

All personnel are highly encouraged to pursue regular GenAI awareness training from multiple sources. This includes foundational knowledge about how GenAI systems function (refer to Appendix Sections A.2-A.4), appropriate use in the workplace, and inherent risks such as hallucinations, bias, and misuse. Departmental resources are provided below, and additional guidance will evolve over time. Personnel are further encouraged to explore the wide variety of public educational content available on the internet.

Resources

The following resources support ongoing learning opportunities:

- DOI Chief Artificial Intelligence Officer (CAIO) (refer to FAQs)
- DOI Artificial Intelligence Training Hub Site

9. Roles and Responsibilities

- **Chief Information Officer (CIO):** Policy co-owner
- **Chief AI Officer (CAIO):** Policy co-owner, risk reviewer, and lead GenAI advocate
- **Department and Bureau Administration:** Ensure staff compliance and training.
- **Enterprise IT and Cybersecurity staff:** Enable, secure, and monitor GenAI tool access.
- **Data Governance staff:** Promotes data quality, transparency, and safeguards
- **All users of GenAI:** Use GenAI responsibly, as described herein.

10. Evaluation and Continuous Improvement

The Department will periodically assess the effectiveness, risks, and benefits of GenAI use through user feedback, performance metrics, and incident reviews. These evaluations will inform updates to policy, training, and tool access.

11. Authorities and References

1. OMB Memorandum M-25-21 -Accelerating Federal Use of AI through Innovation, Governance, and Public Trust (April 3, 2025). https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through- Innovation-Governance-and-Public-Trust.pdf
2. Secretarial Order 3444 – Leading Interior's Path to Artificial Intelligence Transformation, U.S. Department of the Interior (Sept. 17, 2025). https://www.doi.gov/document-library/secretary-order/so-3444-leading-interiors-path-artificial-intelligence
3. DOI Artificial Intelligence Strategy (2025) https://www.doi.gov/sites/default/files/documents/2025-09/doi-ai-strategy.pdf
4. DOI Artificial Intelligence Compliance Plan (2025) https://www.doi.gov/sites/default/files/documents/2025-09/doi-ai-compliance-final.pdf
5. DOI Rules of Behavior for Computer Network Users Reference Guide. https://www.doi.gov/sites/doi.gov/files/uploads/DOI-RULES-OF-BEHAVIOR-FOR-COMPUTER-NETWORK-USERS-REFERENCE-GUIDE.pdf
6. NIST Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1) (July 26, 2024). https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

12. Effective Date

This policy is effective upon the date of signature.

cc: Information Technology Leadership Team
Office of the Solicitor, General Law
Office of Civil Rights
Bureau Chief Financial Officers
Director, Office of Small and Disadvantaged Business Utilization
Competition Advocate

Appendix

A.1 Definitions

Generative AI (GenAI): AI systems capable of generating “new” text, images, code, or other content derived from learned patterns in training data Generated content is considered “new” when it combines or interpolates learned patterns in novel ways, rather than copying directly from training data. However, a human should take care to review the output for accuracy and suitability for its intended use.

Large Language Model (LLM): The engine of GenAI tools, an LLM is a computer-based model trained to recognize and generate human language by identifying patterns from large volumes of data. They may reside in cloud or local environments, but unlike internet search engines that retrieve information from web pages in real-time, LLMs generate responses dynamically from what they’ve learned through *pretraining*. Since pretraining is a computationally expensive process, it is only performed periodically to update LLM knowledge. However, queries to LLMs can be enhanced by adding customized knowledge sources or by using models with real-time web search capabilities. Several well-known LLMs include ChatGPT, Claude, Gemini, and Grok. LLMs are made available as standalone tools or packaged within integrated GenAI tools such as Microsoft Copilot or Amazon Bedrock.

A.2 Understanding Model Behavior

A fundamental concept of GenAI is its ability to generate *probabilistic* output, where the same input query can generate different output variations. While this may seem unsettling, it is the key design feature of GenAI that simulates the variability, adaptability, and creativity in natural human language. By contrast, humans are accustomed to *deterministic* output from computer algorithms, which are predictable and repeatable for a given input. While some GenAI operational modes allow users to control output variability, this does not necessarily improve factual accuracy. The risk of *hallucination*—where the model fabricates output that seems plausible but is actually false—is always a concern. Ultimately, our ability to adapt to output variability and uncertainty from GenAI is the key to discovering its potential value. This requires skilled human judgement to understand the interpretability and trust levels needed for its safe and effective use.

A.3 Use Categories

GenAI tools can be used in a variety of contexts across the Department. This section outlines two broad categories of use— general purpose and special purpose. The examples provided are intended to stimulate experimentation and innovation across different LLMs. Regardless of use category, human oversight and review are required practice for all GenAI outputs that are subsequently incorporated into official documents or materials.

A.3.1 General Purpose Uses

These are typically “out-of-box" capabilities that do not require specialized developer skills. They are assistive and human-in-the-loop in nature and tend to support routine and low-risk applications. Examples include, but are not limited to:

- Drafting and editing documents of all kinds— memos, emails, research papers, reports, summaries, slide content, etc., with human review prior to official use.
- Generating multimedia content— figures, graphics, images, audio, video, etc., to support communication and training.
- Generating computer code or scripts across multiple languages for prototyping, analysis, and automation.
- Assisting with literature reviews, exploratory analysis, and simplifying explanations for complicated subject matter.
- Supporting internal knowledge management (e.g., FAQs, help content, and templates).

A.3.2 Special Purpose Uses

These are typically customized solutions that require some level of specialized developer skills. They are more likely to involve development environments for customization, and/or application programming interface (API) modes with consumption-based costing that provide users with more control over LLM configuration parameters. Agentic AI applications are an example of customized solutions to carry out a sequence of tasks.

A.4 Use Cost Considerations

GenAI tools process data using *tokens*, which are small units of text, image pixel patches, audio segments, etc. Computational costs are monetized through token usage. Each interaction with a GenAI tool consumes input and output tokens, where output tokens are generally more expensive to provide. Tools use a *context window* to throttle or limit token throughput to the LLM. GenAI consumer costs can vary significantly among different provider cost models. For example, fixed cost versus consumption-based cost models will influence tool selection, application, and usage frequency differently. When using consumption-based models, efficient prompt design and user-throttled outputs are an effective way to balance cost with performance needs. Understanding token behavior using consumption monitoring dashboards can provide an effective way to manage costs. Users should expect cost models to evolve with technological advances over time.