

Approved: 12-29-2025

SUBJECT: USE OF GENERATIVE ARTIFICIAL INTELLIGENCE

PURPOSE AND SCOPE

The Department of Energy (DOE or Department) is committed to responsibly harnessing Generative Artificial Intelligence (GenAI) to advance the Department's mission. GenAI refers to AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content, also referred to as “output,” including text, code, images, videos, audio, and other digital media.

This policy sets forth the principles and requirements governing the secure responsible identification, adoption, development, deployment, and use of GenAI tools, both externally and internally developed, recognizing their profound potential to enhance effectiveness, optimize public services, boost internal efficiency, and promote more cost-effective operations. Simultaneously, this policy addresses the inherent risks associated with GenAI, providing guidance to ensure safe and effective adoption of GenAI throughout the Department. This policy empowers DOE employees to strategically leverage GenAI in support of mission-aligned objectives. Existing procurement and security processes remain in effect for AI-related activities. This policy introduces enhanced requirements specifically for the use of GenAI, which are to be implemented in conjunction with, and not as a replacement for, current regulations.

This policy governs all Departmental Elements and their staff at DOE sites (e.g., national laboratories, site offices, and technology centers) involved in using GenAI tools on behalf of the Department. For DOE contractors, the use of GenAI tools, systems, and services remains governed by existing policies including, but not limited to, Federal Information Security Modernization Act (FISMA), Office of Management and Budget (OMB) directives, Executive Orders and standing DOE Orders. Unless otherwise stated, the policy, as implemented by DOE programs, does not apply to recipients of financial assistance from DOE or laboratory-funded technology transfer agreements. GenAI systems must be governed by all existing DOE orders and policies, along with the enhancements required for AI systems.

While the principles and requirements outlined in this policy apply broadly, the Department recognizes that the dynamic environment of scientific research and the pursuit of cutting-edge GenAI innovations may present unique considerations. The Policy is not intended to impose restrictions on foundational GenAI research, experimental development, or the rapid iteration cycles essential for scientific progress.

Therefore, this policy does not apply to the use of GenAI to carry out basic research or applied research, except where the purpose of such research is to develop particular GenAI applications for the Department's or our Nation's use, such as the Genesis Mission. This policy does not cover GenAI when it is being used as a component of a national security system. Additionally, high-impact pilot programs with limited scope and duration, not involving sensitive data or critical operations, may proceed with certification by the Department's Chief AI Officer (CAIO) through an expedited review process. For such internal high-impact pilot programs, individuals interacting with the GenAI should be informed of their participation and have clear avenues for feedback.

In accordance with the responsibilities and authorities assigned by Executive Order 12344, codified at 50 United States Code (U.S.C) sections 2406 and 2511 and to ensure consistency through the joint Navy/DOE Naval Nuclear Propulsion Program, the Deputy Administrator for Naval Reactors (Director) will implement and oversee requirements and practices pertaining to this Directive for activities under the Director's cognizance, as deemed appropriate.

POLICY

The principles and requirements for GenAI use within the Department encompass governance, operational requirements for use, and high-impact use cases. These are outlined as follows:

A. GOVERNANCE

The Department's CAIO, as designated by the Secretary, is responsible for AI innovation, adoption, and governance, in coordination with appropriate agency officials within DOE. The CAIO's duties are defined by OMB. The CAIO will be supported by the Office of Strategy and Technology Roadmaps (OSTR) and the Office of the Chief Information Officer (OCIO) to comply with all duties delineated by current Executive action and direction from the Secretary. The AI Governance Board (AIGB) consists of agency leadership as defined in its charter, is chaired by the Deputy Secretary, and is vice-chaired by the CAIO. The AIGB is responsible for coordinating and governing issues related to AI across the Department.

Individual Departmental Elements (DE) are responsible for assessing their AI needs and available resources. Leadership is required to enable accountable agency officials at the lowest appropriate level to identify, assess, mitigate, and accept risk for AI use cases to more quickly deploy GenAI solutions throughout the Department.

The Chief Information Security Officers (CISOs) and/or the Authorizing Officials (AOs) of individual DEs, Sites, and Labs are responsible for ensuring FISMA compliance for the integration of all GenAI products. They must also ensure Federal Risk and Authorization Management Program (FedRAMP) compliance via the Authority to Operate (ATO) process. Unless otherwise exempted, the integration of GenAI products shall follow the requirements of

existing DOE Directives (DOE O 205.1D). For open-source or internally developed GenAI models not subject to FedRAMP, the CISO or AO must approve an equivalent rigorous security assessment and risk acceptance process in lieu of a standard ATO.

B. OPERATIONAL REQUIREMENTS FOR USE

Human Accountability and Oversight

Human oversight is paramount in the use of GenAI tools. These systems are intended to augment, not replace, human decision-making within the Department.

Responsibility: Users bear ultimate responsibility for the accuracy, completeness, and appropriateness of all GenAI-generated outputs used or relied upon for DOE work. Personnel must be aware that GenAI tools may produce inaccurate, biased, false, or controlled information.

Labeling and Transparency: This policy strongly favors transparency in the use of GenAI for the creation of Department materials. For example, direct outputs of GenAI with a legal or binding effect must be explicitly labeled as AI-generated content. The Department shall fulfill all executive directives on GenAI transparency and disclosure, how GenAI tools operate, their specific objectives, and any other pertinent limitations.

Mandatory Validation and Verification: Users must perform validation checks of GenAI outputs commensurate with the risks and intended use of the output. This process includes:

- **Fact-Checking:** Verifying factual accuracy and logical coherence.
- **Contextual Assessment:** Confirming relevance, appropriateness, and fit within the intended context.
- **Adverse Effect Mitigation:** Critically evaluating outputs for potential inaccuracies originating from hallucinations, exercising diligence to minimize such risks, especially for content influencing decisions or perceptions.

Data

Input Awareness: Personnel shall maintain vigilance and ensure each prompt, whether independent or aggregated, including any restriction on data used in a prompt, adheres to the specific data input restrictions of the GenAI model being used. For example, if the GenAI tool is only approved for open data, data with restrictions greater than open data such as proprietary, Controlled Unclassified Information, or personally identifiable information, may not be used as an input.

Permitted Data Inputs: In the use of GenAI tools, treat all data possessed by DOE in accordance with existing DOE Orders for generating, storing, processing, protecting, and disposing of data. This includes ensuring that all DOE federal employees and contractors adhere

to and follow the requirements of all data markings and security controls (e.g., encryption) for any data used to train a GenAI model, used as a prompt for a GenAI model, or used to create a product output by a GenAI model. DEs, Labs, and Sites must incorporate those requirements that are applicable to their operations.

Unexpected Data Outputs: Spillage of controlled data stemming from GenAI outputs must follow existing DOE data spillage procedures.

Classification Consultation: If any uncertainty exists regarding the classification status or sensitivity level of information intended for input into a GenAI model, personnel must consult with a Derivative Classifier, the designated Classification Officer, or their element's classification subject matter expert prior to any interaction with the GenAI tool. If previously stated officials are not available, consult with your CAIO and/or your CISO.

Data Provenance and Lineage: Compliance with existing DOE and Federal data provenance and lineage requirements is mandatory for GenAI use. Users are accountable for thoroughly documenting all data sources and transformations, especially when GenAI tools generate synthetic data or augment existing datasets. GenAI model data cut-off dates must be considered, particularly for analyses requiring up-to-date information.

Access Control: Controls shall be implemented to manage authorized access to and training/fine tuning of GenAI models, commensurate with the sensitivity of the data involved and to manage any information input, maintained, or output by GenAI models, in accordance with Federal records management requirements.

Intellectual Property

DOE and users of GenAI tools bear primary responsibility for understanding and adhering to all applicable intellectual property (IP) restrictions, including copyright restrictions. These protections and policies will be applied to GenAI as they would to any other activity or technology use within the Department. DOE federal employees are bound by and shall protect confidential information, including adherence to the requirements of 18 U.S.C. §1905 and 15 U.S.C. § 3710a(c)(7) when using GenAI. DOE contractors are bound by and shall adhere to the requirements of the Department of Energy Acquisition Regulation (DEAR) 952.209-72(b)(2)(D) when using GenAI.

Risk of Output Infringement: Many GenAI models are trained on vast datasets that may include copyrighted third-party material. Consequently, users must be aware that outputs generated by these tools could potentially infringe upon existing intellectual property rights, particularly copyrights and trade secrets. Users, working with the service provider, must ensure agreements for GenAI tools address this risk.

Generated Content: Inputs and outputs may fall under the definition of federal/agency records. If they do, such records would be subject to the Freedom of Information Act, 5 U.S.C. § 552, as

amended and Title 10 Code of Federal Regulations Section 1004, 10 C.F.R. § 1004, and may be subject to public release. Ownership of content generated by DOE employees using GenAI tools continues to be governed by 37 C.F.R. Part 501, Uniform Patent Policy for Rights in Inventions Made by Government Employees, and 17 U.S.C. § 105(a), Subject matter of copyright: United States Government works Ownership of content generated by DOE contractors shall be as specified in DOE Acquisition Guide Chapter 70.2701, Patent and Data Rights. Specific guidance on the demarcation between official and personal use will be issued by the OCIO with consultation from the Office of General Counsel.

Training and Awareness

Training Requirements: Requirements for mandatory training for all personnel using GenAI tools within DOE shall be established within six months of the release of this policy. Training shall encompass policy compliance, responsible and ethical GenAI use, mitigating bias and discrimination, identifying GenAI hallucinations (incorrect outputs, false positive outputs, etc.), understanding GenAI limitations and risks, data handling and security protocols when using GenAI, recognizing and reporting potential disinformation or malicious outputs, reporting of security incidents related to the use of GenAI.

Best Practices Sharing: Dissemination of best practices for responsible and secure GenAI use will be conducted primarily through the AI Hub maintained by OCIO and other designated Departmental platforms.

C. HIGH-IMPACT USE CASES

Systems owners are responsible for identifying systems that are or are likely to host high-impact AI use cases as defined by current OMB guidance. While high-impact use cases may be permitted, their deployment is contingent upon the implementation of additional safeguards beyond the best practices established in this policy. It is mandatory for system owners to receive approval from the CAIO prior to the deployment of any high-impact GenAI tool or use case in addition to existing IT procurement policies. This requirement extends to applying previously approved tools to new high-impact use cases.

The CAIO will provide essential guidance for complying with current directives from OMB. The CAIO may issue a waiver exempting a use case from specific OMB requirements, provided such waiver is reported to OMB, publicly posted in a redacted summary protecting national security and sensitive information, and regularly reviewed. Such exceptions will be carefully considered to ensure a judicious balance between accelerating scientific discovery and maintaining robust oversight. All granted exceptions are subject to audit by the CAIO according to federal best practices.


No high-impact AI use case may be deployed without the final approval from the CAIO. This approval signifies that all requirements have been met, ensuring the secure and compliant

deployment of the tool. Any unanticipated incidents with a legal or binding significant effect must be reported to the CAIO. The requirements for high-impact GenAI will take effect starting on April 2, 2026, and any high-impact use case deployed before that date must be brought into compliance or discontinued at that time.

CONTACT

DOE AI Governance Team - DL-AIGovernance@hq.doe.gov

BY ORDER OF THE SECRETARY OF ENERGY:


Secretary of Energy