



ACQUISITION LETTER

This Acquisition Letter is issued under the authority of the Senior Procurement Executives of DOE and NNSA. It is intended for use by the procurement professionals of DOE, primarily Contracting Officers, and other officials of DOE that participate in the acquisition process. Other parties are welcome to its information, but definitive interpretation of its effect on DOE solicitations, contracts, Awards, and other related procedures and actions may only be made by DOE Contracting Officers.

Subject: Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles - Implementation of OMB Memo M-26-04

References:

Executive Orders:

- [Executive Order 14319, Preventing Woke AI in the Federal Government, July 23, 2025](#)
- Executive Order 14179, *Removing Barriers to American Leadership in Artificial Intelligence*, January 31, 2025
- Executive Order 13960, *Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government*, December 3, 2020

OMB Memorandums:

- [Office of Management and Budget Memo M-26-04, Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles, December 11, 2025](#)
- Office of Management and Budget Memo M-25-22, *Driving Efficient Acquisition of Artificial Intelligence in Government*, April 3, 2025
- Office of Management and Budget Memo M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*, April 3, 2025

DOE Policy:

- [DOE P 203.1, Use of Generative Artificial Intelligence](#)

Other References:

- [NIST SP 800-59 – NIST, Guideline for Identifying an Information System as a National Security System](#)

When is this Acquisition Letter (AL) effective?

This AL is effective immediately upon issuance.

When does this AL expire?

This AL remains in effect until superseded or canceled.

Who is the intended audience?

The U.S. Department of Energy (DOE) and National Nuclear Security Administration (NNSA) Contracting Officers (COs). Any reference in this guidance to the DOE should be understood to include the NNSA, unless otherwise indicated.

Who is the point of contact?

DOE questions regarding this AL should be sent to the Contract and Financial Assistance Policy Division at DOE_OAPMPolicy@hq.doe.gov.

NNSA questions regarding this AL should be sent to the Acquisition Policy and Oversight Division, Policy & Oversight Branch (NA-PAS-111), at (505) 845-5639.

What is the purpose of this AL?

This AL provides DOE guidance and required implementation actions in response to Office of Management and Budget (OMB) Memorandum M-26-04, “Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles”. The AL establishes mandatory Department-wide procurement policy, internal controls, documentation standards, and contractual safeguards governing the acquisition, development, integration, and operational use of Artificial Intelligence (AI) systems and services.

What types of contracts are affected by this AL?

This AL is applicable to all solicitations, new contracts and contract modifications where the utilization of AI or robotic/automation processes and/or support is a predominant function of the goods or services under consideration or being acquired. Due to the complexities of applying this policy to basic and applied research and development (R&D), including requirements for processes that determine which kinds of AI models are applicable to this policy, management and operating (M&O) contracts are currently not included in this scope of this AL, and further guidance will be forthcoming.

What is the background information?

On January 31, 2025, President Trump’s administration issued Executive Order (EO) 14179, *Removing Barriers to American Leadership in Artificial Intelligence*. Under this EO, the administration directed agencies to accelerate the federal acquisition and use of AI. Additionally, the EO directed OMB to provide further guidance for implementation of the policies stated in the

order. Subsequently, in compliance with the EO, on April 3, 2025, OMB simultaneously issued M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust* and M-25-22, *Driving Efficient Acquisition of Artificial Intelligence in Government*. Both memos rescinded and replaced prior memos issued by the previous administration related to artificial intelligence acquisition by the federal government.

A follow-up OMB memo, M-26-04, *Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles*, issued on December 11, 2025, further implements EO 14319 to enhance public trust in AI. This memorandum establishes two Unbiased AI Principles: truth-seeking and ideological neutrality, which federal agencies must enforce in their procurement of large language models (LLMs). Agencies are required to include these principles in all new solicitations and orders for LLMs, ensuring compliance through contractual obligations and updated procurement policies. The memo aims to ensure that AI systems used in government remain neutral and nonpartisan, thereby maintaining public trust and institutional credibility.

What guidance is included in this AL?

The following guidance is provided to implement the requirements of M-26-04 within DOE. COs, program offices, and acquisition officials will incorporate those required elements into acquisition planning, solicitation development, evaluation, award, contract administration, and contract close-out for all covered AI procurements.

Sections:

- A. Applicability
- B. Definitions
- C. Implementation Guidance
- D. Additional Information and Guidance

A. Applicability

This Acquisition Letter applies to all Department of Energy (DOE) and National Nuclear Security Administration (NNSA) solicitations, new contracts, and contract modifications where the acquisition of an AI system or service, or the use of Large Language Models (LLMs), is a predominant function of the goods or services under consideration or being acquired. This includes LLMs procured by DOE, regardless of deployment method.

This AL does not apply to AI when it is being used as a component of a National Security System (NSS) as defined in 44 U.S.C. § 3552(b)(6).¹ However, the application of this AL to national security systems, to the extent practicable, is encouraged.

Due to the complexities of applying this policy to basic and applied R&D, including requirements for

¹ AI innovation and risk for national security systems must be managed appropriately, but these systems are governed through other policy. Agencies should reference existing guidelines in place, such as the Department of War's (DoW) Responsible Artificial Intelligence Strategy and Implementation Pathway and the Office of the Director of National Intelligence's Principles of Artificial Intelligence Ethics for the Intelligence Community, as well as policies governing specific high-risk national security applications of AI, such as DoD Directive 3000.09, *Autonomy in Weapon Systems*, <https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodd/300009p.pdf>, and any relevant successor guidance.

processes that determine which kinds of AI models are applicable to this policy, M&O contracts are currently not included in this scope of this AL, and further guidance will be forthcoming.

B. Definitions

For the purposes of this AL, an “AI system or service” aligns with the definition in OMB M-25-21 (e.g., a system or service that uses dynamic or static machine learning algorithms or other forms of AI). Note that widely available commercial products in which AI is embedded as a minor element may be exempt under the “common commercial product” exception, subject to the assessment criteria provided in the memo. An example of such an exception is the Microsoft Office suite of products (i.e. Word, Excel, etc.). For further information on exemptions, see Section 2 of OMB M-25-21.

For the purposes of this AL, the following additional definitions apply:

- **AI (Artificial Intelligence):** The term “artificial intelligence” has the meaning provided in Section 238(g) of the John S. McCain National Defense Authorization Act for Fiscal Year 2019 (NDAA of 2019). This encompasses machine learning (including deep learning), reinforcement learning, transfer learning, and generative AI, but generally excludes robotic process automation or systems defined solely by human-defined rules. The NDAA of 2019 states that AI is a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. AI systems use machine learning and other techniques to learn how to achieve objectives. For procurement purposes, this includes machine learning, deep learning, neural networks, and generative AI, but excludes simple automation or rule-based systems.
- **AI Impact Assessment (AIIA):** a documented, lifecycle-based evaluation conducted prior to contract award (and updated as required) that identifies, analyzes, and mitigates the operational, ethical, legal, cybersecurity, privacy, workforce, and mission impacts of acquiring, deploying, or using an artificial intelligence capability on behalf of the Federal Government, to inform acquisition decisions, risk acceptance, and required approvals
- **AI Model:** A component of an information system that implements AI technology and uses computational, statistical, or machine-learning techniques to produce outputs from a given set of inputs.
- **AI System:** A data system, software, application, tool, or utility that operates using dynamic or static machine learning algorithms or other forms of artificial intelligence. It includes systems primarily for AI research, development, or implementation, or where AI capability is integrated into other systems or business processes. It generally excludes common commercial products with embedded AI, such as word processors or map navigation systems, that have substantial non-AI purposes or functionalities.
- **Agency:** As defined in 44 U.S.C. § 3502(1), any executive department, military department, Government corporation, Government controlled corporation, or other establishment in the executive branch of the Government, including independent regulatory agencies. It does not include the Government Accountability Office, Federal Election Commission, or government-owned contractor-operated facilities engaged in national defense research and

production activities.

- **Bias (Algorithmic Bias):** Systematic prejudice or unfairness in AI outcomes. In federal procurements, this term carries significant weight due to the government's obligation to serve all citizens equitably. Procurement specifications often explicitly require bias detection, mitigation, and reporting mechanisms.
- **Chief AI Officer (CAIO):** An official designated by the head of each agency within 60 days of the issuance of OMB M-25-21 to promote AI innovation, adoption, and governance. For CFO Act agencies, the CAIO must hold a Senior Executive Service (SES) or equivalent position.
- **CFO Act Agencies:** Agencies identified in 31 U.S.C. § 901(b). The Department of Energy is a CFO Act agency.
- **Generative AI (GenAI):** AI systems capable of generating text, code, images, audio, video, or synthetic data.
- **High-Impact AI:** AI with an output serving as a principal basis for decisions or actions that have a legal, material, binding, or significant effect on civil rights, civil liberties, privacy, access to essential services (education, housing, insurance, credit, employment, critical government resources), human health and safety, critical infrastructure, public safety, or strategic assets and resources. See also OMB Memo M-25-21.
- **Human-in-the-Loop:** A governance mechanism requiring human review prior to finalization of AI-assisted decisions.
- **Large Language Models (LLMs):** A generative AI model trained on vast, diverse datasets that enable the model to generate natural-language responses to user prompts. Large, typically transformer-based language models trained on massive text corpora to generate sequences of tokens (by predicting the next token), enabling them to produce unstructured text (natural language) and structured text outputs (e.g., code) in response to prompts.
- **Machine Learning (ML):** (Consistent with general definition) A subset of AI enabling systems to learn from data. For federal procurements, the types of data (classified, PII, public-sector data) and their handling become particularly critical.
- **National Security System:** A National Security System is an information system designated under the Committee on National Security Systems (CNSS) policy and federal guidelines as essential to national security missions, requiring specialized classification handling, enhanced cybersecurity protections, and heightened supply-chain oversight. It includes systems used for intelligence, military, nuclear security, and other operations where compromise could cause exceptionally grave damage to national security. Such systems must comply with CNSS requirements, NIST SP 800-59, and DOE/NNSA directives governing classification, safeguarding, cryptography, procurement, and risk management. As defined in 44 U.S.C. § 3552(b)(6).
- **OMB M-26-04 Compliance:** Refers to the requirements outlined in OMB Memorandum M-26-04, *Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles*,

specifically the “truth-seeking” and “ideological neutrality” principles for LLMs.

- **Open Government Data:** Data collected or produced by the federal government that is made available to the public in an open, machine-readable format. Federal AI procurements often leverage or are required to produce insights from such data, promoting transparency and public benefit.
- **Responsible AI:** A framework and set of principles that guide the development, deployment, and use of AI systems to ensure they are fair, accountable, transparent, robust, and beneficial to society. This is a paramount concern in federal AI procurement, often mandated by Executive Orders (e.g., EO 13960, “Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government”) and agency policies.
- **Trustworthy AI:** Often used interchangeably with Responsible AI in federal discourse. It emphasizes attributes such as explainability, fairness, robustness, safety, privacy-enhancement, and security. Federal procurements frequently require solutions to adhere to Trustworthy AI principles.
- **Unbiased AI Principles:** The two principles identified in E.O. 14319 and further elaborated in OMB M-26-04:
 - Truth-seeking: LLMs will be truthful in responding to user prompts seeking factual information or analysis, prioritizing historical accuracy, scientific inquiry, and objectivity, and acknowledging uncertainty where information is incomplete or contradictory.
 - Ideological Neutrality: LLMs will be neutral, nonpartisan tools that do not manipulate responses in favor of ideological dogmas, and developers will not intentionally encode partisan or ideological judgments into outputs unless prompted or readily accessible to the end user.

C. Implementation Guidance

DOE will not acquire, develop, deploy, or authorize the operational use of AI systems or services without documented compliance with this AL. All AI acquisitions must incorporate risk-based safeguards proportionate to the system’s impact level.

1. Market Research and Planning

COs will determine whether the requirement includes AI functionality during acquisition planning. This determination must be documented in the Acquisition Plan or contract file. When AI functionality is present, the requiring activity/office will initiate an AI Impact Assessment (AIIA).

2. AI Impact Assessment (AIIA)

Before issuing a solicitation (or before award for simplified acquisitions), the requiring activity/office, with coordination with the OCIO and the CAIO, will complete an AIIA.

AIIA process should be tailored based on impact categorization focusing on high impact systems while reducing the admin burden for low and moderate impact systems. The assessment must address the following:

- Impact categorization (Low, Moderate, High Impact)
- Civil rights and privacy implications
- Cybersecurity considerations
- Data governance and data quality
- Bias testing requirements
- Explainability requirements
- Human oversight design
- Lifecycle monitoring and controls

High Impact AI systems require written concurrence from the CAIO, legal review, Privacy Officer review (as applicable), and OCIO cybersecurity approval prior to contract award.

3. Performance-Based Acquisition and Solicitation Requirements

COs and Contracting Officer's Representatives (CORs) should ensure the following is included in solicitations and award requirements documents (i.e., a performance work statement (PWS) or statement of work (SOW), etc.):

- a) Clear statement of intended AI use case and operational boundaries.
- b) Minimum accuracy and performance thresholds.
- c) Mandatory bias testing and mitigation plan submission.
- d) Model documentation and transparency disclosures.
- e) Data provenance and data governance description.
- f) Security architecture and compliance certifications.
- g) Lifecycle management and update notification procedures.
- h) Evaluation criteria addressing responsible AI and risk mitigation capability.

4. Cybersecurity and Information Security Requirements

AI systems need to comply with DOE cybersecurity directives, FedRAMP (where applicable), NIST AI Risk Management Framework guidance, secure software development

lifecycle practices, records management laws and procedures, and continuous monitoring requirements. Contractors are required to maintain audit logs, protect against adversarial manipulation, and implement secure data pipelines. For more information on these topics, please reach out to the DOE/NNSA Office of the Chief Information Security Officer.

5. Testing, Validation, and Continuous Monitoring

Prior to operational deployment, AI systems must undergo validation testing for accuracy, bias, security vulnerabilities, and operational robustness. Post-deployment monitoring includes periodic reporting, performance reassessment, and disclosure of material model updates.

6. Evaluation, Award, and Contract Management

When soliciting proposals and awarding contracts, clearly define evaluation and performance criteria to uphold Unbiased AI Principles. This will result in a contract requirement tied to eligibility, payment terms, and grounds for termination for non-compliance if violated. Vendors must provide basic transparency documentation for each LLM purchase, supporting regular reporting, updated assessments, notifications of model updates, and independent testing—such as prompt-pair bias analysis, ambiguity resolution, and drift monitoring—to ensure and verify ongoing compliance. For deployments with greater risk, implement expanded transparency requirements that address activities before and after training, thorough model evaluations, enterprise-level oversight, and insight into third-party alterations.

Requirements for New Solicitations, Contracts and Modifications:

Effective immediately, COs will ensure that any solicitation or order for the procurement of an AI system or service which utilizes LLM, or any AI model other than an LLM to which the agency determines the requirements of OMB M-26-04 should apply, including contractual requirements addressing compliance with the Unbiased AI Principles. These requirements, as outlined in OMB M-26-04, are considered material to contract eligibility and payment. Non-compliance may lead to contract termination. COs will incorporate the attached contract clause, “**Ensuring Unbiased AI Principles (APR 2026)**,” into all applicable solicitations and contracts.

Requirements for Existing Contracts:

To the extent practicable, COs should modify existing contracts for AI systems or services with LLM modeling to include the attached contract clause, “**Ensuring Unbiased AI Principles (APR 2026)**,” addressing compliance with the Unbiased AI Principles. This modification should occur no later than the exercise of any option that extends the period of performance of the contract.

For AI models other than LLMs in existing contracts, agencies should consider the relevant factors identified in Appendix A.2 of OMB M-26-04 to determine whether to apply these requirements.

7. AI Inventory and Reporting

Program Offices are required to register all AI systems in the DOE AI Inventory and update records annually or upon material system change. High-Impact AI systems will be reported in accordance with OMB reporting instructions.

D. Additional Information and Guidance

- Agencies must maximize the use of American AI by prioritizing American-developed AI tools and services.
- Agencies must protect privacy by establishing policies and processes to ensure compliance with privacy requirements in law and policy, especially when dealing with Personally Identifiable Information (PII).
- Agencies must preserve government data and intellectual property (IP) rights by including appropriate contractual terms that clearly delineate ownership and IP rights. Contracts should protect government data from unauthorized disclosure or use and prevent its use for training vendor commercial offerings without explicit agency permission.

Attachment A to AL 2026-05 Contract Clause

Prescription: The Contracting Officer will insert the following clause in solicitations and contracts for technical and management support services, advisory and assistance services, architectural and engineering services, and information technology services involving AI systems. This clause flows down to all subcontractors at all tiers. Therefore, the Contractor will insert a clause, modified appropriately to substitute “Contractor Representatives” for “the Contracting Officer” in all subcontracts. This clause does not apply to Management and Operating (M&O) contracts.

Ensuring Unbiased AI Principles (APR 2026)

(a) Unbiased AI Principles - General Obligation

(1) Unbiased AI Principles.

The Contractor will design, configure, operate, and maintain all Artificial Intelligence (AI) Systems and Large Language Models (LLMs) provided under this Contract in accordance with the following Unbiased AI Principles:

- (i) truth-seeking, meaning the AI System is oriented toward generating factual, grounded, and accurate outputs with documented limitations; and
- (ii) ideological neutrality, meaning the AI System will not systematically favor or disfavor any political party, ideology, or viewpoint in its default behavior and will function as a neutral tool serving the Government and the public on a nonpartisan basis.

(2) Human accountability.

The Contractor acknowledges that AI outputs are advisory only and will not be represented as replacing human judgment for Government decisions carrying legal or significant effects; the Contractor will not design or configure the AI System in a manner that obscures or misrepresents the limits of the system’s reliability or authority.

(b) Truth-Seeking and Accuracy

(1) Factuality and grounding.

The Contractor warrants that the AI System has been developed and configured using reasonably designed processes to promote factual, grounded outputs, including use of curated data sources, evaluation on appropriate benchmarks, and implementation of safeguards to mitigate hallucinated or fabricated content.

(2) Disclosure of limitations.

The Contractor will document known limitations and failure modes of the AI System, including domains in which outputs are less reliable, and will provide clear end-user guidance and warnings describing such limitations and appropriate human review.

(3) Continuous improvement.

Upon notice from the Government of repeated or systemic inaccurate or misleading outputs in Government-relevant use cases, Contractor shall, at no additional cost, implement reasonable corrective measures, which may include configuration changes, additional training or tuning, and updated documentation, within a mutually agreed timeframe.

(c) Ideological Neutrality and Non-Discrimination

(1) Neutral default behavior.

The Contractor will ensure that, absent a user's explicit request for a particular perspective, the AI System's default responses do not promote or disadvantage any political party, ideology, or viewpoint, and do not intentionally embed partisan framing in generic informational queries.

(2) Testing for bias.

The Contractor will conduct and periodically update documented evaluations reasonably designed to detect systematic ideological, demographic, or viewpoint bias in outputs for Contractor provided, government approved Government-defined test suites, including "prompt-pair" and ambiguity tests where feasible, and will provide summary results to the Government upon request.

(3) Remediation of identified bias.

Where the Government identifies material instances of systematic bias inconsistent with this clause, Contractor will promptly propose and implement remediation steps, including model or configuration adjustments and updated safeguards, at no additional cost to the Government.

(d) Transparency Artifacts

(1) Documentation.

The Contractor will provide and maintain, for the term of this Contract and at no additional cost, the following documentation for each covered AI System:

- (i) An acceptable use policy applicable to Government and end users.
- (ii) Documentation of capabilities, intended uses, limitations, data characteristics where reasonably knowable, and known risks.
- (iii) A data card or equivalent Documentation describing data handling practices, privacy protections, and relevant security controls.
- (iv) End-user guidance and instructions to support responsible, truth-seeking, and neutral use.

(2) Incident and issue reporting.

The Contractor will provide mechanisms for users to report problematic outputs (including biased, harmful, or factually incorrect content) and will maintain a process for triaging, prioritizing, and addressing such reports, providing periodic summaries to the Government on request.

(3) Change notification.

The Contractor will provide advance written notice of material changes to model architecture, training data sources (to the extent known and disclosable), safety policies, or content-moderation configurations that are reasonably likely to affect factuality, neutrality, or safety of outputs used by the Government.

(e) Materiality, Remedies, and Enforcement

(1) Material requirements.

The parties agree that compliance with the Unbiased AI Principles and the transparency obligations in this Contract are material requirements of performance and will be conditions of the Government's obligation to accept deliverables and make associated payments.

(2) Cure and re-performance.

If the Government reasonably determines that Contractor has materially failed to comply with this clause, the Government may issue a written notice of non-compliance. Contractor shall, within a specified cure period, submit and implement an acceptable corrective action plan, including re-performance, additional testing, mitigation steps, and updated documentation, at no additional cost.

(3) Withholding and other remedies.

If Contractor fails to cure non-compliance within the specified period, the Government may, in addition to any other rights and remedies: (i) withhold payment on affected deliverables, (ii) suspend use of specific AI functionalities, (iii) reflect non-compliance in past-performance evaluations, and (iv) pursue any other remedy available at law or under this Contract, including termination as provided elsewhere.

(f) Enhanced Transparency

(1) Pre- and post-training information.

At the Government's request, Contractor will provide information sufficient to describe:

- (i) system prompts and safety policies that materially affect model behavior,
- (ii) pre-deployment red-teaming and adversarial testing performed, and

- (iii) post-deployment monitoring programs used to detect bias, drift, or degradation in accuracy.

(2) Model evaluations.

The contractor shall, on an annual basis or as otherwise specified by the Government, re-verify that the AI system remains in compliance with the principles of this clause. The Contractor will maintain and provide, upon request, evaluation documentation describing methods and outcomes for: (i) bias and fairness testing, (ii) benchmarks related to accuracy and reliability in domains relevant to the Government's use, and (iii) tests of ambiguous versus straightforward prompts for ideological neutrality. The Contractor will provide a summary report of the re-verification results to the Contracting Officer upon request.

(3) Enterprise-level controls.

The Contractor will provide features or configuration options that enable the Government to:

- (i) manage access and roles,
- (ii) log and audit interactions consistent with applicable law and policy, and
- (iii) monitor model performance and safety safeguards over time.

(4) Third-party modifications.

If Contractor provides a derivative or integrated model based on a third-party model, Contractor will disclose the nature of its modifications, fine-tuning, or additional training, and explain how such modifications are reasonably expected to affect behavior, bias, or risk profile in a report deliverable upon the government's request.

(End of clause)