



In Bid to Ban “Woke AI,” White House Imposes Transparency Requirements on Contractors

What You Need to Know

Key takeaway #1

On December 11, 2025, the Office of Management and Budget issued guidance meant to bar the federal government from purchasing AI models that incorporate “ideological biases or social agendas.” While the guidance does not prohibit the purchase of any particular AI, it requires contractors to provide agency customers with information on how AI models are built, trained, or modified.

Key takeaway #2

The guidance has several notable exceptions, such as excluding AI models used on national security computer systems, open-source models, and models used incidentally by contractors for administrative purposes.

Key takeaway #3

Vendors should carefully review new contract clauses to ensure compliance and reduce the risk of contract termination or legal exposure.

Client Alert | 7 min read | 12.19.25

In July 2025, President Trump signed Executive Order (EO) 14319, *Preventing Woke AI in the Federal Government*, to preclude the federal government from procuring artificial intelligence (AI) models that incorporate “ideological biases or social agendas,” including “diversity, equity, and inclusion.” The EO mandates that the federal government purchase only large language models (LLMs) developed according to two “Unbiased AI Principles” — that they be “truth-seeking” and show “ideological neutrality.” To implement these principles, the EO directed the Office of Management and Budget (OMB) to issue guidance.

On December 11, 2025, the OMB published **Memorandum M-26-04** to implement the ban on ideological, untruthful LLMs (“guidance”). The guidance will impose transparency, documentation, and disclosure obligations on AI contractors, or “vendors.”

Because the U.S. federal government is the **world's largest purchaser** of goods and services and its strictures can impact how products are developed, the guidance amounts to the federal government's most significant regulation yet of AI.

However, the guidance is relatively limited. It does not prohibit the purchasing of any particular LLM, instead only requiring that the vendor provide the agency with information about how the LLM is built, trained, and modified. The guidance also has several notable exceptions, such as excluding "national security systems" managed by the Department of Defense; LLMs acquired pursuant to free, open-source licenses; and LLMs used incidentally by a contractor for administrative purposes, among others.

By March 11, 2026, all agencies must revise procurement policies consistent with this guidance. Contractors will need to review new contract clauses carefully to ensure compliance.

"Unbiased AI Principles"

The guidance operates in parallel with OMB's earlier memoranda governing AI acquisition and oversight, including **EO 13960, *Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government***, OMB Memoranda **M-25-22, *Driving Efficient Acquisition of Artificial Intelligence***, and **M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust***. The guidance seeks to implement the "Unbiased AI Principles" that the EO identifies as "truth-seeking" and based on "ideological neutrality":

- *Truth-Seeking*: LLM outputs must be "truthful in responding to user prompts seeking factual information or analysis," prioritize historical and scientific accuracy, and "acknowledge uncertainty" when information is incomplete or contradictory.
- *Ideological Neutrality*: LLMs must be "neutral, nonpartisan tools," and developers must not "intentionally encode partisan or ideological judgments into an LLM's outputs," unless the user so prompts.

New Transparency Requirements

The guidance requires that all agencies revise procurement policies to include contractual requirements for transparency as well as remedies for potential noncompliance. For what the guidance terms "minimum" and "enhanced" transparency requirements (detailed below), agencies should identify which requirements are "material to eligibility and payment under the contract, to support termination of the contract for default . . . in cases where the vendor refuses to take corrective action in identified cases of noncompliance."

OMB also directs agencies, "to the extent practicable," to modify existing contracts for LLMs to include such requirements and implement them prior to exercising any contract options. Updated policies must also include processes for LLM users to report outputs that violate the "Unbiased AI Principles." The guidance expires on December 11, 2027.

Minimum Threshold for LLM Transparency: The guidance requires agencies to request from vendors:

- *An acceptable use policy* that differentiates between appropriate and inappropriate uses of the LLM.
- *Model, system, and/or data cards*, which include "summaries of the training process, identified risks and mitigations, and model evaluation scores on LLM benchmarks."

- *End-user resources*, which may include product tutorials, developer guides, or other tools to help users “maximize utility” of the LLM.
- *An end-user feedback mechanism*, such as an email inbox or point of contact, so users can report violative outputs.

Threshold for Enhanced LLM Transparency: An agency may choose to request information beyond the minimum requirements. Such additional information may focus on:

- *Pre- and post-training activities* that would impact the “factuality and grounding of LLM outputs,” such as guidelines on responding to user-generated queries. They may also include information on the types of outputs restricted by “content moderation and safety filters,” data on the use of red teaming “to protect against incidents of bias in generated output,” training or development that occurred outside the United States, and modifications made to comply with regulations from other governmental bodies.
- *Model evaluations*, including the results of and methodology for bias evaluations, as well as benchmark scores for vendor evaluations to measure a model’s bias, among others.
- *Enterprise-level controls*, including governance and model evaluation tools, along with features that require models to cite sources for outputs.
- *Third-party modifications*, namely disclosure of additional controls to modify LLMs’ outputs, if the vendor is not a direct LLM developer.

Scope

Far-reaching: The guidance applies to “any LLM procured by an agency, regardless of the manner in which the LLM will be deployed, further modified, or used by the agency.” It also directs agencies to request information regarding LLM development and operation when those models are “integrated” into software or services that the agency is procuring. For models developed by an agency itself, the guidance encourages the agency to integrate these principles by documenting the model’s pre-training, evaluation, and controls and including a user-feedback mechanism.

Exceptions: While far-reaching in what LLMs it covers, the guidance includes many significant carveouts. The guidance does not apply to “**national security systems**,” LLMs “acquired pursuant to a free, open-source license,” and AI used “incidentally” by a contractor during performance of a contract for “administrative purposes,” among others. (In the absence of compulsory requirements, the guidance encourages agencies to establish procedures to perform due diligence on the alignment of open-source models with the EO and OMB Memorandum M-25-21.)

The guidance's implementing appendix also suggests agencies impose the requirements flexibly, depending on the circumstances of the vendor. For example, while agencies must obtain sufficient information from the vendor to determine whether the LLM "complies with the Unbiased AI Principles," the guidance recognizes that the information available "will vary depending on the vendor's role within the software supply chain," and more information will be available to some vendors than others. Similarly, the guidance recognizes that agencies often transact with third-party LLM providers that will not have access to all of the LLM's product information. The guidance directs agencies to "consider these nuances of LLM procurement in determining how to apply" the requirements. The guidance instructs agencies to "avoid" imposing requirements on a vendor to disclose sensitive technical data, such as model weights, which are the numerical parameters in a neural network that determine how it processes inputs to produce an output.

Finally, for models other than LLMs, including models that generate images or audio, the guidance should "inform the documentation requirements imposed for the procurement" only "where practicable." This flexibility is notable, considering that the EO specifically criticized AI models for "chang[ing] the race or sex of historical figures . . . when prompted for images."

Impact

The guidance does not forbid any particular LLM. It only requires agencies to demand information on how models are built, tested, and deployed. Many AI developers are already providing at least some of this information to the public. Nevertheless, vendors should take steps to comply, particularly because of the short compliance window and because failure to fulfill "material" requirements could be grounds for termination or lead to risks of bid protests, liability under the False Claims Act, or suspension/debarment.

Contractors can consider:

- Assessing whether existing products meet the "Unbiased AI Principles";
- Inventorying existing model documentation (e.g., model cards, data cards, system prompts, evaluation reports);
- Reviewing pre- and post-training activities, model evaluations, enterprise-level controls, and third-party modifications, and including a mechanism for user feedback;
- Evaluating reseller and integrator agreements to ensure compliance with these new policies and ensuring flowdown clauses to capture subcontractors and AI component providers; and
- Training on the obligations of this and related White House guidance.

Crowell & Moring LLP and Crowell Global Advisors will continue to monitor U.S. Government efforts to adopt, promote, and regulate AI.

For further information, please contact our team.

Contacts

Matthew F. Ferraro

Partner

Washington, D.C. D | +1.202.624.2610

mferraro@crowell.com

Neda M. Shaheen

Associate

She/Her/Hers

Washington, D.C. D | +1.202.624.2642

nshaheen@crowell.com

Kate M. Crowley

Partner, Crowell Global Advisors Senior Director

Washington, D.C. D | +1.202.624.2698

Washington, D.C. (CGA) D | +1.202.624.2500

kgrowley@crowell.com

Michael G. Gruden

Partner

Washington, D.C. D | +1.202.624.2545

mgruden@crowell.com