

Occupational Safety and Health Review Commission



Artificial Intelligence Compliance Plan for OMB Memorandum M-25-21

Date: 08/22/2025

Chief AI Officer, Senior Attorney Advisor: Ron Bailey
Telephone Number: (202) 606-5410
E-mail Address: CAIO@oshrc.gov

I. About the Review Commission

The Occupational Safety and Health Review Commission (Review Commission) is a small, independent federal agency that plays a vital role in ensuring safe and healthy working conditions for American workers. The Review Commission provides fair and timely adjudication of work-related safety and health disputes between employers, employees or their representatives, and the Occupational Safety and Health Administration (OSHA), an agency of the United States Department of Labor. OSHA establishes workplace safety and health standards and inspects workplaces to enforce compliance with those standards.

The Review Commission was created by the Occupational Safety and Health Act of 1970 as an independent agency that is entirely separate from OSHA and the Department of Labor. This separation ensures that all parties involved in an enforcement proceeding—OSHA, employers, affected employees, and authorized employee representatives—receive an impartial hearing before the Review Commission.

The cases that come before the Review Commission arise from citations that OSHA issues employers following workplace inspections. The Review Commission's adjudicative process begins when an employer files a notice contesting the OSHA citation or an employee or employee representative files a notice contesting the abatement date stated in the OSHA citation. Following receipt of both the notice of contest and citation, the Review Commission's Office of the Executive Secretary assigns the case a docket number, creates a new case file, and notifies all parties of the case's docketing. Next, the Review Commission's Chief Administrative Law Judge assigns the case to one of the Review Commission's administrative law judges (ALJs) for disposition. Parties may seek discretionary review of an ALJ's decision from the agency's three-member, presidentially appointed Commission.

II. Driving AI Innovation

A. Removing Barriers to the Responsible Use of AI.

Requests to use AI are evaluated on a case-by-case basis, taking into consideration (1) the availability of the Review Commission's budgetary and human resources; (2) the nature of the requested AI use case, including how it relates to and would affect the agency's adjudicatory mission; and (3) whether the requested AI use case, if implemented, would comply with the requirements of OMB Memoranda M-25-21 and M-25-22.

The Review Commission's use of AI is limited by its adjudicatory function and limited resources. However, the agency recognizes that as AI applications advance over time, their output will likely become more accurate, and the costs associated with use of AI may become more manageable. The Review Commission, therefore, will assess on a continuing basis whether it is appropriate to incorporate any requested AI use cases into the adjudicatory process.

B. Sharing and Reuse

The Review Commission's AI use cases are from commercial-off-the-shelf products or services that incorporate AI-based tools (for example, Westlaw Edge or certain Microsoft products). The Review Commission does not use AI code, models, and data sets beyond what is incorporated into these products or services. Given the Review

Commission's limited resources, as well as its adjudicatory function, the agency has no plans to develop AI applications unique to the agency.

C. AI Talent

To the extent budgetary resources are available, the Chief Artificial Intelligence Officer (CAIO) and the Chief Information Officer (CIO) will actively pursue AI training opportunities to enhance their understanding of AI-related issues at the Review Commission. Any requests for AI training from other employees will be evaluated by the CAIO, in consultation with the requesting employee's supervisor and the Executive Director, on a case-by-case basis, taking into consideration how the training relates to the employee's work at the Review Commission and the availability of budgetary resources.

The approval of any employee training is subject to the requirements included in the Review Commission's internal directive on Employee Development and Training.

III. Improving AI Governance

A. Agency Policies

By December 29, 2025, the Review Commission will review and, if necessary, revise internal agency directives that concern privacy and computer and information security to ensure that the directives comply with the requirements of M-25-21. The Review Commission's directive on procurement and contracting will also be reviewed and, if necessary, revised to ensure compliance with the requirements of both M-25-21 and M-25-22.

The Review Commission plans to develop an internal policy for the use of generative AI by December 29, 2025, the deadline identified in M-25-21 for developing such a policy. This policy will set the terms for acceptable use of generative AI at the Review Commission and establish adequate safeguards and oversight mechanisms that allow generative AI to be used by the agency, when appropriate, without posing undue risk.

B. AI Use Case Inventory

Annual information request. Annually, the CAIO sends all agency employees an email that (1) describes what types of AI uses may constitute AI use cases for purposes of M-25-21; (2) notifies them that the use of AI is permissible as long as it complies with current policies, privacy regulations and laws, and any other relevant legal requirements; and (3) directs them to provide a description of any AI uses to the CAIO through a Microsoft Form. Once information on the agency's use of AI is collected, the CAIO—in collaboration with the Senior Agency Official for Privacy (SAOP) and the CIO—will determine which uses must be included in the agency's AI Use Case Inventory. This process also allows the CAIO, SAOP, and CIO to assess whether the agency's uses of AI are in accordance with current policies and legal requirements.

Continuous updates. The AI Use Case Inventory will be updated by the CAIO on a continuing basis, including when (1) an additional use of AI is reported by agency personnel or is otherwise evident (for example, from a new technology being acquired through the agency's procurement process); and (2) discontinuation of an AI use is reported

by agency personnel or is otherwise evident (for example, from cancellation of a subscription service).

IV. Fostering Public Trust in Federal Use of AI

According to M-25-21, “AI is considered high-impact when its output serves as a principal basis for decisions or actions that have a legal, material, binding, or significant effect on rights or safety.” More specifically, M-25-21 defines “high-impact AI” as follows:

AI with an output that serves as a principal basis for decisions or actions with legal, material, binding, or significant effect on:

1. an individual or entity’s civil rights, civil liberties, or privacy; or
2. an individual or entity’s access to education, housing, insurance, credit, employment, and other programs;
3. an individual or entity’s access to critical government resources or services;
4. human health and safety;
5. critical infrastructure or public safety; or
6. strategic assets or resources, including high-value property and information marked as sensitive or classified by the Federal Government.

Section 6 of M-25-21 lists categories for which the expected use of AI that serves as a principal basis for agency decision or action is presumed to be high-impact.

After collecting information on all the Review Commission’s AI use cases (see Part III.B), the CAIO, in collaboration with the SAOP and CIO, will evaluate each use case to determine whether it should be characterized as “high-impact AI” under M-25-21. In determining whether a use is high impact, the following factors will be considered: (1) the AI’s specific output; (2) the extent to which that output serves as a *principal* basis for decisions or actions that have a legal, material, binding, or significant effect on one or more of the categories specified in the above definition; and (3) the potential risks to the rights or safety issues specified in the above categories.

The Review Commission currently has no high-impact AI use cases and does not plan to pursue uses of AI that would qualify as “high impact” under M-25-21. Therefore, the agency has not developed (1) agency-specific minimum risk management practices for high-impact AI, or (2) agency-specific procedures for issuing, denying, revoking, certifying, and tracking waivers for one or more of the minimum risk management practices. However, to ensure compliance with the requirements of section M-25-21, the Review Commission will develop and implement such practices and procedures by April 3, 2026.

Given that the Review Commission does not plan to pursue any use of AI that would qualify as high-impact, the agency is presently in compliance with M-25-21’s requirement to safely discontinue any high-impact AI use case not compliant with the requirements in section 4 of the memorandum.