



UNITED STATES OFFICE OF PERSONNEL MANAGEMENT
Washington, DC 20415

The Director

MEMORANDUM

TO: Heads of Departments and Agencies
FROM: Scott Kupor, Director, U.S. Office of Personnel Management
DATE: August 27, 2026
RE: Use of Artificial Intelligence in the Federal Hiring Process

I. Summary

Agencies may be compromising efficiency and quality in Federal hiring by failing to adopt the use of Artificial Intelligence (AI). Executive Order 14170 and the *Merit Hiring Plan* require agencies to improve Federal hiring by developing better position descriptions and assessments, increasing efficiency in hiring, and reducing average time-to-hire to 80 days.¹ AI is an important tool to help implement these goals and should be leveraged in concert with requirements outlined in OMB M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*.²

To align policy and Federal hiring goals, this memorandum³ provides guidance on how agencies may approach hiring AI use cases, and advises agencies to adopt AI across the hiring lifecycle under the conditions below.⁴

¹ Executive Order 14179, [Removing Barriers to American Leadership in Artificial Intelligence](#) (Jan. 23, 2025); Executive Order 14170, [Reforming the Federal Hiring Process and Restoring Merit to Government Service](#) (Jan. 20, 2025); White House Domestic Policy Council & OPM, [Merit Hiring Plan](#) (May 29, 2025); Chance to Compete Act of 2024, Pub. L. No. 118-188.

² Nothing in this memorandum is intended to impose requirements that are more restrictive than, or that otherwise modify, the requirements of OMB M-25-21. Rather, it is intended to clarify the interpretation and application of M-25-21 in the Federal hiring context and to support agency adoption of AI. In the event of any inconsistency between this guidance and OMB M-25-21, M-25-21 controls.

³ This memorandum applies to the internal management of the Executive Branch and does not create any right or benefit, substantive or procedural, enforceable at law or in equity by any party against the United States, its departments, agencies, or entities, its officers, employees, or agents, or any other person.

⁴ OMB Memorandum M-25-21, [Accelerating Federal Use of AI through Innovation, Governance, and Public Trust](#) (Apr. 3, 2025). Citations are to the numbered sections of the M-25-21 Appendix. Footnote 27 states that AI “may be integrated in decision or activity pipelines in high-impact categories without meeting the definition

II. High-Impact Standard and Analysis

In accordance with M-25-21, agencies should treat a use of AI as “high impact” if the AI’s “output 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; (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.” Thus, agencies should look to 1) whether AI outputs serve as a principal basis for a decision or action; and 2) whether the decision or action will have a legal, material, binding, or significant effect on one or more of the interests enumerated in the definition above, including, as relevant in the Federal hiring context, an individual or entity’s access to employment.

In OPM’s view, where an appropriate agency official independently reviews supporting data, evaluates the decisional logic, and adopts the AI-assisted work product, the output generally does not serve as a principal basis for the resulting decision or action. That is not because human oversight, by itself, exempts a use from the high-impact definition. Instead, it is because independent review, evaluation, and adoption supply the decision’s own justification. Agencies should not treat this analysis as a barrier to adoption; Table 1 in Appendix A applies this standard to common hiring use cases.

Some categories of AI use cases are “presumed high-impact” as described in Section 6 of OMB M-25-21, including: “[d]etermination of the terms or conditions of Federal employment, including pre-employment screening, reasonable accommodation, pay or promotion, performance management, hiring or termination, or recommending disciplinary action; reassignment of workers to new tasks or teams.”⁵

For AI use cases that fall under this and other presumed high-impact categories, agencies must treat the use case as high-impact, unless they determine that, despite the presumption, the use does not actually fit within OMB’s definition of the term. If an agency makes such a determination, an “appropriate agency official must submit written documentation to notify the CAIO [Chief Artificial Intelligence Officer]” of that fact. For AI use cases that remain high-impact, agencies must either:

of high-impact because the AI’s output does not actually ‘serve as a principal basis for’ the relevant type of agency action or decision.”

⁵ Agencies taking a careful read of M-25-21 will note that not all AI use cases for hiring fall under this presumption.

1. Implement minimum risk management practices; or
2. Submit a waiver request to the agency CAIO for one or more minimum risk management requirements (an example of a waiver justification is provided in *Appendix B – Example Model Determination Language*).

III. Assessing Common Uses of AI in Federal Hiring

To provide more guidance on how to evaluate common uses of AI for hiring as being high-impact or not, OPM has analyzed several common use cases. Please reference “Table 1. Approaches to Assessing Common Uses of AI in Federal Hiring” in *Appendix A* of this Memorandum to see a practical framework for evaluating common AI uses in Federal hiring. The uses in the table are illustrative, not exhaustive, and agencies are encouraged to pursue other beneficial uses of AI across the hiring lifecycle, consistent with OMB M-25-21. The assessments below and in Appendix A reflect OPM’s views on the application of OMB’s guidance to common hiring fact patterns, developed in consultation with OMB; they are not official OMB determinations, and agencies remain responsible for making their own determinations under M-25-21. After following the assessment process outlined in Appendix A, OPM considers the following uses of AI in Federal hiring to be generally not high-impact under M-25-21:

1. **Content development for job opportunity announcements:** this includes drafting the package of artifacts developed for a new job opportunity, including position descriptions, evaluation statements, announcement text, job analyses, qualification requirements, crediting plans, and interview questions, when the process of drafting content itself does not affect an individual’s access to employment. Before operational use, the material must be reviewed, validated, and adopted under applicable personnel requirements (classification under 5 U.S.C. Ch. 51; job analysis and validation under 5 CFR 300.103; subject matter expert (SME) involvement), obligations that apply independently of M-25-21.
2. **Decision-support for applicant-level evaluations:** this includes résumé screening, minimum qualifications and eligibility review, and candidate summarization when the AI output is not a principal basis for a decision or action. For example, when the AI locates, extracts, and summarizes the applicant’s documentation and how it aligns with agency-adopted eligibility and qualification criteria, with a reference back to the underlying documentation, and the deciding official independently reviews the underlying record and makes the determination on that basis, the AI output generally will not serve as a principal basis for the decision, because the official’s independent review of the record — rather than the AI summary — supplies the decision’s justification.

3. **Pre-offer quality control:** this includes review of the staffing case file and selection before a final job offer, when the AI is limited to administrative completeness and quality-control support and is not a principal basis for approving, delaying, or rejecting the selection or final offer.
4. **Aggregate hiring program evaluation:** this includes analysis of completed hiring actions, applicant flow, time-to-hire, and assessment performance, where no individual employment decision turns on the AI output.

Agencies must still assess their specific workflow and deployment context to affirmatively reach a non-high-impact determination, and an appropriate agency official must document that determination to the CAIO.

A non-high-impact determination changes nothing else. AI-assisted steps remain subject to job analysis and validation requirements under 5 CFR 300.103. Agencies must verify, before deployment and periodically, that:

1. AI-assisted rating, ranking, and referral correctly apply veterans' preference;
2. AI-enabled steps meet Rehabilitation Act accessibility and reasonable accommodation requirements;
3. Applicant data is processed only in agency-authorized tools with applicable Privacy Act and privacy impact assessment protections;
4. Each hiring action remains documented and auditable; and
5. A reconsideration avenue remains available for adverse eligibility determinations.⁶

Agencies should aim to streamline the processes where possible. Agencies are reminded that a high-impact designation is not a prohibition on adoption. Chief Human Capital Officers (CHCOs) are encouraged to collaborate with their CAIOs and other relevant stakeholders (e.g., Chief Privacy Officers, Chief Information Officers, Chief Data Officers, etc.) to implement the required and recommended actions in Table 1, including the minimum risk management practices of OMB M-25-21 § 4(b). This is especially true as many policies still apply to the overarching use case, regardless of whether an AI use case is high-impact under M-25-21.

⁶ This item creates no new process or right; it restates avenues that exist independently of this memorandum — for example, correction procedures under agency delegated examining authority and veterans' preference redress under 5 U.S.C. § 3330a — and requires only that AI-assisted steps not impede access to them.

IV. What is Next?

OPM will continue introducing new AI features across the USA Suite of tools, including USA Staffing, USAJOBS, USA Hire, USA Class, and USA Performance. However, agencies should not wait for those features to begin adopting AI consistent with this memorandum.

OPM will also address AI-assisted processes in delegated examining oversight, and assist agencies on request, including through the Artificial Intelligence Governance and Adoption Program (AI GAP). Questions regarding this memorandum may be directed to AIGAP@opm.gov.

cc: Chief Human Capital Officers, Deputy CHCOs, Human Resources Directors, Chief Information Officers, Chief Artificial Intelligence Officers, and Chiefs of Staff

Appendix A

Appendix B

Appendix A – Table 1. Approaches to Assessing Common Uses of AI in Federal Hiring

The assessments in this table reflect OPM’s views on the application of OMB M-25-21 to common Federal hiring use cases, developed in consultation with OMB. Agencies remain responsible for their own determinations under M-25-21.

AI Use Case Outputs	STEP 1: Is the Use Case Presumed High-Impact or does it meet the High-Impact Definition?	STEP 2: Is the AI being used as a “Principal Basis”? Assess whether the AI’s outputs are used for a decision or action that will have a legal, material, binding, or significant effect.	STEP 3: Are there other relevant considerations?
<i>Identify the Common Uses of AI in Federal Hiring</i>	<i>Determine whether this AI use applies to any of the categories for which the use or expected use of AI that serves as a principal basis for an agency decision or action is presumed to be high-impact (e.g., Determination of the terms or conditions of Federal employment, including pre-employment screening, reasonable accommodation, pay or promotion, performance management, hiring or termination, or recommending disciplinary action; reassignment of workers to new tasks or teams). The presumption may be overcome only where a documented factual basis establishes that the use does not meet the definition of high-impact AI (see M-25-21 App. 14), and an appropriate agency official must document</i>	<i>Is the AI being used as a “principal basis” for decisions or actions that have a legal, material, binding, or significant effect on the interests enumerated in the M-25-21 definition of high-impact AI, including an individual or entity’s access to employment? What are the potential impacts?</i> <i>The inquiry is functional, not formal: identify what in fact supplies the justification for the decision or action in the documented workflow. An AI output may serve as a principal basis even where a human formally approves the result; conversely, an output that merely informs a decision that is independently justified by the record does not.</i> <i>A high-impact determination is possible whether there is or is not human oversight for the decision or action. Asserted human validation or review does not, by itself, remove a use from the definition; what matters is whether the documented record shows that the official’s independent judgment, rather than the AI output, supplied the decision’s justification.</i>	<i>Consider the following when applying this use of AI.</i>

	<i>any such determination to the CAIO.</i>		
Draft position descriptions, evaluation statements, and job announcements	This use case may fall under the presumed high-impact category. However, OPM would typically not consider this specific AI use high-impact because it is usually not a principal basis for a decision or action. The next step is determining if the AI use serves as a principal basis for an agency decision or action.	In OPM’s view, this specific AI use is typically not a principal basis for a decision or action when the draft is a starting point that the responsible official independently reviews, revises, and adopts, ensuring alignment with the classifier’s authority and the OPM Qualification Standards Handbook, because the official’s independent review and adoption — rather than the AI draft — supply the justification for the final material. The adopted material derives its operative force from that verification against governing standards, not from the draft itself. By contrast, where the responsible official adopts AI-drafted material without substantive, documented review, the draft supplies the justification for what is adopted, and the use should be treated as high-impact. Agencies are encouraged to align with OMB M-25-21 when evaluating the AI’s specific output and its potential risks.	Applicable classification law and policy (5 U.S.C. Ch 51) continue to apply. The material must be reviewed and adopted by an appropriate agency official before use.
Draft job analyses, qualification requirements, crediting plans, and interview questions	This use case may fall under the presumed high-impact category. However, OPM would typically not consider this specific AI use high-impact because it is usually not a principal basis for a decision or action.	In OPM’s view, this specific AI use is typically not a principal basis for a decision or action when the draft is a starting point that the responsible officials independently review, revise, and adopt, because their independent review and adoption — rather than the AI draft — supply the justification for the adopted criteria and instruments. Because these instruments later govern how applicants are evaluated, independent validation under 5 CFR 300.103 — against the position’s actual duties and documented job-relatedness — is the safeguard that keeps the AI draft from serving as a	Agency adopts all criteria and instruments before they are used to evaluate applicants (SME involvement; 5 CFR 300.103). The adoption step must include independent validation of the AI-drafted material against the position's actual duties and

	The next step is determining if the AI use serves as a principal basis for an agency decision or action.	principal basis for downstream applicant determinations. Where that validation is not performed and documented before use, the use should be treated as high-impact. Agencies are encouraged to align with OMB M-25-21 when evaluating the AI's specific output and its potential risks.	documented job-relatedness (5 CFR 300.103) before use.
Screen résumés and match against minimum qualifications; review documents to determine veterans' preference; review documents to determine other eligibilities (e.g., reviewing Interagency Career Transition Assistance Plan (ICTAP) documentation, reviewing documentation for Schedule A for individuals with disabilities, reviewing military spouse documentation, selective service	This specific AI use is considered presumed high-impact, as this output may serve as a principal basis for agency decisions or actions that have a legal, material, binding, or significant effect on the interests enumerated in M-25-21's definition of high-impact AI, including access to employment. The presumed list of high-impact AI categories relates to this AI use include pre-employment screening, hiring, and reasonable accommodation. The next step is determining if the AI use serves as a principal basis for an agency decision or action.	In OPM's view, this specific AI use <i>is typically not a principal basis</i> for a decision or action when: - the AI locates, extracts, and summarizes that record against the applicable eligibility criteria; - an authorized official who reviews the documentation — including (but not limited to) the AI's summary of that documentation and the entirety of the package, evaluates it against the governing eligibility requirement, makes the determination on that basis; and - an authorized official justifies the determination by reference to the underlying record, ensuring that the official would reach the same result on review of the source documents themselves and that the official would rely on the entirety of the package including (but not limited to) the summary. The traceability and availability of the documentation is critical to supply independent justification. Where officials in practice ratify AI outputs without independent review of the underlying record, the output serves as a principal basis notwithstanding the nominal review step, and the use must be treated as high-impact with the minimum risk management practices of M-25-21 § 4(b) applied. Agencies are encouraged to align with	The conditions are that outputs identify the specific documentation each determination rests on, that the documentation remains accessible to the deciding official, and that accuracy is established through sampling and quality assurance sufficient to support the official's reliance — with particular attention to determinations that would remove an applicant from consideration.

registration documents, U.S. citizenship, etc.); recommend eligibility determinations		OMB M-25-21 when evaluating the AI's specific output and its potential risks.	
Generate a proposed recommendation of applicants who do not meet minimum qualifications; assess candidates and generate match rationales to support referral and selection decisions	This specific AI use is considered presumed high-impact, as this output may serve as a principal basis for agency decisions or actions that have a legal, material, binding, or significant effect on the interests enumerated in M-25-21's definition of high-impact AI, including access to employment. The presumed list of high-impact AI categories relates to this AI use include pre-employment screening, hiring, and pay or promotion. The next step is determining if the AI use serves as a principal basis for an agency decision or action.	In OPM's view, this specific AI use is typically not a principal basis for a decision or action when: • an authorized official reviews the documentation — including (but not limited to) the AI's summary of that documentation and the entirety of the package and evaluates it against the governing eligibility requirement and makes the determination on that basis; • an authorized official justifies the determination by reference to the underlying record, ensuring that the official would reach the same result on review of the source documents themselves and that the official would rely on the entirety of the package including (but not limited to) the summary; and • the record and the official's judgment supply the justification, ensuring that the rationale is traced to identifiable evidence in the application material. Adverse recommendations warrant particular care: where the official's recorded rationale merely restates the AI-generated rationale without independent verification against the application material, the AI output supplies the justification, serves as a principal basis, and the use should be treated as high-impact. Agencies are encouraged to align with OMB M-25-21	In practice: the rationale must cite the specific application material and job-analysis factors it rests on, and the selecting official's independent conclusion must be recorded in the file. Otherwise, treat as high-impact and apply § 4(b) minimum practices.

		when evaluating the AI's specific output and its potential risks.	
Rate, rank, or score candidates against the adopted crediting plan; AI-scored assessments that produce candidate scores or rank order — with or without review and certification by an appropriate human official (including AI as sole scoring instrument)	This specific AI use is considered presumed high-impact, as this output may serve as a principal basis for agency decisions or actions that have a legal, material, binding, or significant effect on the interests enumerated in M-25-21's definition of high-impact AI, including access to employment. The presumed list of high-impact AI categories relates to this AI use include pre-employment screening and hiring. The next step is determining if the AI use serves as a principal basis for an agency decision or action.	In OPM's view, this specific AI use is typically not a principal basis for a decision or action when: • an authorized official reviews and certifies the crediting plan that the AI supplies and the results — including (but not limited to) the AI's summary of that documentation and the entirety of the package, evaluates it against the governing eligibility requirements, and makes the determination on that basis; • an authorized official justifies the determination by reference to the underlying record, ensuring that the official would reach the same result on review of the source documents themselves and that the official would rely on the entirety of the package including (but not limited to) the summary; • the record and the official's judgment supply the justification, ensuring that the rationale is traced to identifiable evidence in the application material; and • each score traces to identified plan elements and the applicant's own responses such that an authorized official can review the derivation — including in summary form — confirm that the score is the one the adopted plan yields, and certify the result on that basis, the plan and the applicant's submission supply the justification. Where scores are produced by means that cannot be traced to the adopted plan and the applicant's responses, where the derivation is not reviewable by the certifying	

		official, or where AI operates as the sole scoring instrument without a reviewable basis for certification, the score is itself the underpinning of the decision, no independent record supports it, and the output likely serves as a principal basis whether or not a human official certifies the result. In that circumstance, the use should be treated as high-impact and the minimum risk management practices of M-25-21 § 4(b) applied. Agencies are encouraged to align with OMB M-25-21 when evaluating the AI's specific output and its potential risks.	
Review staffing case file and selection before issuance of a final job offer to identify missing documentation, inconsistencies, procedural errors, or potential compliance concerns	This use case may fall under the presumed high-impact category. However, OPM would typically not consider this specific AI use high-impact because it is usually not a principal basis for a decision or action. The next step is determining if the AI use serves as a principal basis for an agency decision or action.	In OPM's view, this specific AI use is typically not a principal basis for a decision or action when the AI output informs an HR professional when a candidate's package may have missing information so that the agency can follow up, any flag that could remove a candidate from consideration is verified against source documentation, and the output does not influence a candidate's ranking, rating, or performance-based eligibility. Where the output would have such an influence, OPM's view is that the use should be treated as high-impact and the minimum risk management practices of M-25-21 § 4(b) applied. The limiting principle is that approving, delaying, or withholding a final offer affects an individual's access to employment; the conditions above accordingly confine the AI's role to surfacing administrative issues for human follow-up, so that any action affecting the candidate rests on verified source documentation rather than on the flag itself.	An authorized HR official reviews the AI output, evaluates and resolves relevant issues, and makes the final determination on whether the selection and final job offer may proceed. Two conditions apply. The absence of a flag may not substitute for a required certification by the authorized official, and any flag that would remove a candidate from consideration must be verified against source documentation before action is taken.

		Agencies are encouraged to align with OMB M-25-21 when evaluating the AI's specific output and its potential risks.	
Evaluate completed hiring actions and recruitment outcomes; analyze applicant flow, time-to-hire, certificate and selection patterns, and assessment performance; produce quality-of-hire and other program-evaluation analytics; assess whether completed hiring actions were carried out consistent with merit system principles; and identify adverse impact or other anomalies in aggregate hiring data to support process improvement and human capital accountability.	This use case may fall under the presumed high-impact category. However, OPM would typically not consider this specific AI use high-impact because it is usually not a principal basis for a decision or action. The next step is determining if the AI use serves as a principal basis for an agency decision or action	In OPM's view, this specific AI use is typically not a principal basis for a decision or action. The outputs are statistical and programmatic: they describe completed hiring actions and aggregate patterns rather than any pending application, so they do not underpin a decision or action with a legal, material, binding, or significant effect on any individual's enumerated interests. No decision affecting an individual issues from the analysis; the output informs process improvement and human capital accountability. If aggregate outputs are later repurposed to inform a decision or action concerning an identifiable individual, that is a distinct use of AI that requires its own assessment under this framework. Agencies are encouraged to align with OMB M-25-21 when evaluating the AI's specific output and its potential risks.	Aggregate analyses should use the minimum applicant data necessary and remain subject to applicable Privacy Act and privacy impact assessment protections. If an analysis suggests an error in an individual, completed hiring action, the matter should be referred through existing correction, audit, or oversight processes rather than acted on directly from the AI output.

Appendix B – Example Model Determination Language

(For use by the appropriate agency official in the written documentation to the CAIO required by OMB M-25-21 § 4(a). Adapt bracketed text and any other aspect as appropriate. This model language is OPM's and is offered for convenience of agency use, not as a requirement. A determination that a presumed high-impact use case does not meet the definition of high-impact (see M-25-21 App. 14) and a request to waive one or more minimum risk management practices (see M-25-21 App. 15) serve different purposes and should be prepared and submitted as separate documents, not combined.)

[Use case name] involves the use of [AI tool] to [match applications against agency-adopted qualification requirements and crediting plans / generate candidate summaries and match rationales /rate, rank, or score candidates against agency-adopted crediting plans or assessments] in support of [hiring action type].

This use case [does/does not] fall under OMB's list of presumed high-impact AI, in accordance with M-25-21.

If presumed high impact but not a principal basis: Although this use case is presumed high-impact, under [Agency]'s implementation [describe the specific facts of the workflow — e.g., the independent review, evaluation, and adoption steps — that establish why the AI output does not serve as a principal basis], it is not a “principal basis” for decisions or actions that have a legal, material, binding, or significant effect on the interests enumerated in M-25-21's definition of high-impact AI. Therefore, it is not considered high-impact. [The presumption may be overcome only where the documented facts establish that the definition is not met.]

If presumed high impact and a principal basis: This use case is presumed high-impact. Under [Agency]'s implementation [insert details], it is a “principal basis” for decisions or actions that have a legal, material, binding, or significant effect on the interests enumerated in M-25-21's definition of high-impact AI. Therefore, it is considered high-impact. [Explain how Agency will meet Minimum Risk Management Practices or a Waiver.]

Example waiver justification (prepare and submit as a separate document where applicable): [Agency] requests that the CAIO waive [identify the specific minimum risk management practice(s) of M-25-21 § 4(b)] for [use case name]. Applying [the practice] to this use case would [increase risks to safety or rights overall / create an unacceptable impediment to critical agency operations] because [explain]. The following compensating controls remain in place: [describe, e.g., human review and certification, sampling and audit logs, performance monitoring]. The waiver is requested for [scope and duration] and will be reassessed [upon material change to the tool, workflow, or deployment context, and at least annually]. This waiver and its justification will be documented, tracked, and reported consistent with M-25-21.

Based on the workflow, controls, testing, and monitoring described above, [name/title] determines that this use:

[] Is high-impact AI because the AI output serves as a principal basis for a consequential employment decision. § 4(b) of M-25-21 applies.

[] Is not high-impact AI because the AI does not serve as a principal basis for an agency action or decision with legal, material, binding, or significant effect on the interests enumerated in M-25-21's definition of high-impact AI.

Basis for determination:

[Insert concise explanation.] This determination will be revisited upon a material change to the tool, criteria, workflow, or context, or if testing or audit results show that the review controls described above are not functioning as documented.