National FOIA Office
U.S. Environmental Protection Agency
1200 Pennsylvania Avenue NW (2310A)
Washington, DC 20460

Re: Appeal of Determination in FOIA Request 2026-EPA-03349

Dear FOIA Appeals Officer:

Introduction

This letter is an appeal of the above-captioned determination issued by the EPA via email on June 23, 2026. Public Employees for Environmental Responsibility ("PEER"), by and through its legal counsel, the Free Information Group, PLLC ("FIG"), appeals the sufficiency of EPA's search for records responsive to PEER's FOIA request, its determination to withhold records under Exemption 5, and its failure to properly segregate information in withheld records.

PEER's request sought two tranches of information regarding EPA's use of Artificial Intelligence ("AI") tools and its compliance with OMB and other legal requirements regarding such tools. For simplicity, this appeal refers to these as parts A and B. Those requests sought:

Part A – Use of AI Tools in Chemical Review

1. Records of prompts used, and refinements to previously-used prompts for the AI screening rules used in risk assessments for the following chemicals, listed by Chemical Abstracts Service (CAS) registry number:

a. 1918-00-9: Dicamba (3,6-dichloro-2-methoxybenzoic acid)

b. 50-00-0: formaldehyde

c. 13188-82-4: 2-Propenoic acid, (2,2-dimethyl-1,3-dioxolan-4-yl)methyl ester

d. 151249-82-0: Methanone, 1,1'-(octylphosphinylidene)bis[1-(2,4,6-trimethylphenyl)-

e. 1955538-52-9: 2-Propenoic acid, 2,2-dimethyl-1,3-dioxan-5-yl ester

f. 2489703-47-9: Benzene, 1-methoxy-4-[3-methyl-4-(2-phenylethoxy)-3-buten-1-yl]-

g. 288859-91-6: Poly(oxy-1,2-ethanediyl), -(2-methyl-1-oxo-2-propenyl)--(tridecyloxy)-, branched

2. The contents and change history of:

a. prompts stored in EPA OCSPP prompt libraries relating to TSCA;

b. system prompts;

c. AI configuration files and data (including model name and ad-hoc fine tuning/RAG scripts or tools);

d. runbooks or prompting guidance for human users of AI tools.

3. Technical guidance memos from the AI Digital Services Consulting Team to OCSPP.

4. Technical consulting reports and read-ahead memos from the AI Digital Services Consulting Team regarding TSCA Systematic Review pilots.

5. Post-Deployment Monitoring records associated with the AI Chemist Assistant, EcoVault, or any other AI tool used in TSCA reviews.

Part B – Compliance Documentation

1. The full version of and metadata for the 2025 AI Use Case Inventory, as reported to OMB;

2. All final impact assessments for high-impact AI tools used by EPA in 2025;

3. Any waiver requests discussed by or submit to the Data Governance Advisory Committee AI Subcommittee;

4. Any notices sent to the OMB CAIO that EPA has determined a particular AI use which would be presumptively high-impact will not be considered high-impact;

5. Documentation of EPA's compliance with high-impact AI requirements in OMB Memorandum M-25-21 §4(b);

6. System prompts or instruction tuning documents used to align EPA AI models with the administration's "Unbiased AI Principles;"

7. Any bias evaluation or benchmark scores relating to the bias of any AI used by EPA and the methodology used for evaluating those scores;

8. Any AI Values Alignment Report prepared for any deployed EPA AI system;

9. Any governance tools, such as customizable system instructions additive to a base model's system prompts or content generation filters relating to any AI used by EPA; and

10. Internal directives or memoranda regarding development of the Generative AI Rules of Behavior, including timelines for completion, staffing, and the review process to be followed in its development.

Appeal

Public Employees for Environmental Responsibility ("PEER"), by and through its legal counsel, the Free Information Group, PLLC ("FIG"), appeals the following determinations:

1. EPA's Office of Chemical Safety and Pollution Prevention ("OCSPP") has no records responsive to Part A(1-5)
2. EPA's Office of Finance and Administration ("OFA") has no records responsive to Part B(3-10), and
3. OFA's assertion that the Final Impact Assessments for high-impact EPA AI tools may be properly withheld under Exemption 5 as protected by the deliberative process privilege.

Insufficient Search – Part A(1-5)

In follow-up communications to its initial determination, EPA has represented that its no-records response is based on a finding that AI is not in use for OCSPP chemical assessments. This is irreconcilable with publicly reported facts and statements by EPA regarding the use of AI in its systematic review process.

EPA has spent the last 18 months announcing the introduction and development of new AI tools across all its scientific programs. PEER's initial request referenced a publication titled *AI tools used in EPA's Systematic Review Process* which, while unavailable in full, states that it summarizes "AI methods currently used in chemical assessments and investigations into how generative AI (GenAI) could be used for even more tasks."[1]

[1] Watford, S., AI TOOLS USED IN EPA'S SYSTEMATIC REVIEW PROCESS (Aug. 12, 2025), https://assessments.epa.gov/risk/document/&deid%3D366912.

In its request, PEER also cited an Amazon Web Services ("AWS") blog post about a presentation delivered by EPA staff regarding the agency's use of AI at the AWS Summit Washington, DC 2025.[2] That blog post stated that EPA's "'Powering the Great American Comeback' initiative includes a pillar to make the U.S. the artificial intelligence (AI) capital of the world. . . . OCSPP and [the] Office of Pesticide Programs (OPP) are putting those words into action." The blog described EPA-developed proof-of-concept platforms developed with AWS to "help EPA scientists to evaluate studies."

The AWS blog explained that an EPA scientist at OCSPP "aimed to increase the rate at which his team could provide high-quality results for chemical assessments" by developing a generative AI tool. That tool would "handle the laborious document analysis" involved in qualitative evaluation of studies based on criteria specific to the goals of particular chemical risk assessments. The blog concluded that the experiment successfully proved that foundation models ("FMs") running on Amazon Bedrock can assist EPA scientists to accelerate mission delivery of chemical reviews, and that it could be improved by "optimizing the study quality criteria language to get better AI-generated outputs."

EPA has developed and is further developing a prompt library and iterative refinements of system prompts to optimize model output.[3] If EPA is working on changing how its AI judges the quality of scientific studies, those instructions exist as system prompts.[4]

---

[2] Dan Ford, Sandra El Ashry, Brian Zambrano, Steve Goley, Natasha Tchir, Newel Hirst, and Sameer Mohamed, *Accelerating workflows with generative AI: EPA's document processing journey*, AWS PUBLIC SECTOR BLOG (July 14, 2025), https://aws.amazon.com/blogs/publicsector/accelerating-workflows-with-generative-ai-epas-document-processing-journey/.

[3] Further supported by the explanation that "Where model inferences were not the same as previous human reviews... it was typically found that the prompt could be improved to facilitate more accurate judgement."

[4] The AWS documentation also refers to "evaluation prompts" fulfilling the same role. The role of system prompts was explained in PEER's request. System prompts are instructions that define an AI model's behavior and operational goals before user interaction begins. These directives are generally persistent, guiding the AI's tone, safety protocols, and output format throughout the entire conversation. They are a reflection of agency policy in code form. OCSPP uses a framework called Population, Exposure, Comparator, and Outcome ("PECO") to formulate a research question and facilitate literature review. AI implements this framework based on criteria defined in its prompts.

EPA has also specifically requested an additional $2.5 million in appropriations for FY27 "to accelerate the use of AI approaches in chemical evaluations."[5]

EPA is so confident in the efficiency and effectiveness of its AI chemical evaluation program, that it has proposed reducing the office's budget by $13 million, or 30.3 total *workyears* of staff capacity, "to increase the efficiency of EPA operations and reduce duplicative, voluntary, and unnecessary work."[6] The agency expects to find efficiencies by "[e]xpanding the use of artificial intelligence in EPA's systematic review process for existing chemical risk evaluations, which will reduce costs while increasing the speed of review and the consistency of results from this foundational work."[7]

PEER thus appeals the agency's determination that no AI is used in the EPA's chemical review process.

PEER similarly appeals the determinations that no system prompts, user guidance, technical memoranda, and other evaluation records from the AI Digital Services Team described in Part A exist. Perhaps the agency just has an issue with the terminology used in PEER's request. So PEER will now clarify so the agency can fulfill its legal obligations to do a sufficient search. If these documents do not exist as files called "system prompts" or "technical guidance memos," the agency is requested to search for and produce the equivalent materials used in the design, operation, and evaluation of the above-described systematic review AI project. Put simply, PEER's request seeks the agency's instructional inputs that direct how the documented Amazon Bedrock model evaluates scientific literature and the policy guidance that informs the agency's AI deployments.

As AWS explained, "EPA has reported its intent to use AI to assist staff . . . . for scientific work" and emphasized the importance of "ensur[ing] scientific integrity of the results." PEER agrees with the foundational importance of scientific integrity. As such,

---

[5] E.P.A., FY 2027 BUDGET OVERVIEW, TAB 00 – OVERVIEW, V (April 2026), *available at* https://www.epa.gov/planandbudget/cj; *see also id.*, at VII; IX-X ("To modernize its operations and improve efficiency, EPA will significantly invest in expanding its internal AI capabilities. EPA will make strategic adjustments to its programs and activities to incorporate the use of AI across internal and external processes and systems."); *id.*, at TAB 02 GOAL AND OBJECTIVE OVERVIEW, p. 15; Tab 03, p. 116 ("Continued modernization of artificial intelligence (AI)-enhanced tools, models, and databases . . . will help mitigate potential delays in chemical evaluation, resulting in faster decisions under TSCA.").

[6] *Id.*, at 350.

[7] *Id.*, at 351.

PEER cannot accept EPA's assertion that its well documented AI pilot program simply does not exist.

Insufficient Search Part B(1)

EPA replied to Part B(1) of PEER's FOIA request with "the public 2025 AI Use Case Inventory."[8] PEER specifically requested the "full version of and metadata for the 2025 AI Use Case Inventory, as reported to OMB."

PEER is concerned that, like other agencies, EPA has publicly underreported its AI use case inventory. On Aug. 13, 2026, FEDSCOOP reported on an appearance by Katie Noyes, the Chief AI Officer for the Federal Bureau of Investigation, at the Department of Defense Intelligence Information System (DoDIIS) Worldwide Conference 2026.[9] In her appearance, CAIO Noyes stated "[w]e've got, by the way, 139 approved use cases — all different manners of data, all different manners of processes and certainly all manner of risk." The FBI's reported AI use case inventory for 2025 included just 50 use cases. PEER's request also noted that the OMB template for use-case reporting appears to contain more data fields than included in EPA's public use case inventory. While the 34 columns of data are greater than the three columns available at the time of PEER's request, it still lags the more than 50 fields present in the 2024 OMB template.

PEER is also concerned that the provided use case inventory indexes Use Case ID numbers up to 73, but lists only 19 rows of information.

PEER requested that EPA provide the full version of the use case inventory as reported to OMB. PEER appeals the sufficiency of EPA's search which produced only the "public" version of the document. If the produced document was the final version as reported to OMB, EPA should state as much. If it was not, EPA must produce the version provided to OMB or claim an exemption to FOIA justifying its withholding.

Insufficient Search – Part B(3-10)

EPA's public statements about its high-profile AI investments are well documented in the initial request and above. Since the request was filed, EPA has further developed the record of its "development and refinement of AI use cases with high operational impact" and how they "will contribute to innovative solutions for pressing

[8] Determination Letter from Janice Jablonski to Kevin Bell (June 23, 2026) (filename "2026-EPA-03349 FD Letter.signed.pdf").
[9] Lindsey Wilkinson, *FBI is using much more AI than what it previously disclosed*, FEDSCOOP (Aug. 13, 2026), https://fedscoop.com/fbi-ai-use-case-disclosure-caio-katie-noyes/.

environmental challenges" and "yield significant achievements."[10] In addition to the TSCA AI projects described above, EPA has explained to Congress that it "is developing robust AI capabilities that propel the United States to the forefront of global AI leadership, while ensuring operational efficiency."[11] These activities fall under the agency's stated operational goal: "Goal: Make America the AI Capital of the World." EPA has explained one such project is even "a pilot project using AI tools in Agency rulemaking."[12]

EPA's no records response is unsupported by the agency's public statements about its own compliance obligations and the realities of AI deployments. EPA is publicly boasting about its many ongoing and planned high impact AI uses. Meanwhile, it has represented in response to PEER's FOIA request that it has no documentation of *any* compliance with OMB M-25-21, beyond, implicitly, EPA's now 11-month-old compliance plan itself.[13] EPA represents it has no governance tools or system instructions for any of its dozens of reported use cases or its unknown quantity of undisclosed use cases. EPA's "no records" response indicates that the agency has taken no steps to align its AI models with relevant executive orders, and made no efforts to evaluate the bias of AI models used in its work, including the project to have AI write EPA's own regulations. The agency also apparently has no documented plan to develop the "Generative AI Rules of Behavior," nor has it assigned staff to the project. If this is truly the case, EPA seems profoundly uninterested in its stated goals of making America the AI capital of the world, no idea how or why its models do anything, and no desire to learn.

It is far more likely that EPA simply hasn't looked very hard for responsive documents. PEER appeals EPA's refusal to conduct a good faith, legally sufficient search for the technical and compliance material requested. This documentation exists, and must exist, according to EPA's own policies on the subject. PEER appeals EPA's refusal to conduct a search reasonably designed to uncover all material relevant to Part B(3-10) of PEER's FOIA request.

---

[10] E.P.A., FY 2027 BUDGET OVERVIEW, TAB 00 – OVERVIEW, X (April 2026), *available at* https://www.epa.gov/planandbudget/cj

[11] *Id.*, at TAB 04 – ENVIRONMENTAL PROGRAMS AND MANAGEMENT 277.

[12] *Id.*, at 293.

[13] EPA Compliance Plan for OMB Memorandum M-25-21 (Oct. 2025), https://www.epa.gov/data/ai-compliance-plan.

Exemption 5: Deliberative Process Privilege – Part B(2)

Finally, PEER appeals the withholding of EPA's final impact assessments for high-impact AI tools used by EPA in 2025 under the deliberative process privilege. The agency's stated reasoning for the withholding is merely that "they are predecisional and deliberative and would harm agency decision making if released[.]" This cursory denial, barely a line of text stating the legal standard absent any reasoning for its application, is manifestly insufficient to properly invoke the privilege or to establish harm to any interest it may protect.

The deliberative process privilege protects material that is generated prior to a decision in order to inform that decision, and which is so sensitive or individual that public exposure would harm the agency's decision-making process. It has several purposes, including "to assure that subordinates within an agency will feel free to provide the decisionmaker with their uninhibited opinions and recommendations without fear of later being subject to public ridicule or criticism; to protect against premature disclosure of proposed policies before they have been finally formulated or adopted; and to protect against confusing the issues and misleading the public by dissemination of documents suggesting reasons and rationales for a course of action which were not in fact the ultimate reasons for the agency's action."[14]

Records which are unrelated to or come at the end of the give-and-take of the policymaking process are not deliberative and cannot be withheld under the privilege. Records whose release would not foreseeably harm the agency's policy making process may not be withheld under Exemption 5. EPA "must concretely explain how disclosure 'would'—not 'could'—adversely impair internal deliberations."[15]

It is unclear what is deliberative about these records, as the documents themselves are designated "final" by the terms of PEER's FOIA request, and "final" documents or policies cannot be withheld under deliberative process privilege. These documents contain no candid expressions of opinion regarding delicate policy decisions. As an *ex post* assessment of an agency IT platform conducted solely for compliance purposes, it is likely not even predecisional to the adoption of the AI tool in question. EPA's released inventory references multiple use cases which have been implemented while the assessment is still pending. EPA has not described which or even how many records are

---

[14] *Coastal States Gas Corp. v. Dep't of Energy*, 617 F.2d 854, 866 (D.C. Cir. 1980)

[15] *Reps. Comm. for Freedom of the Press v. Fed. Bureau of Investigation*, 3 F.4th 350, 369–70 (D.C. Cir. 2021) (quoting *Machado Amadis v. United States Dep't of State*, 971 F.3d 364, 371 (D.C. Cir. 2020)).

.

withheld. This is precisely the kind of boilerplate denial-by-default approach which led Congress to conclude "Exemption five has been singled out as a particularly problematic exemption. Some have taken to calling it the 'withhold it because you want to' exemption."[16]

There is also no plausible harm to agency candor or decision making from release of these final technical reports. There is no employee candor which may be reduced, no process which may be harmed by the release of this information, except that it may expose a cursory and unserious process for the approval of dangerous AI tools and lead to public pressure to slow down their adoption. The only harms foreseeable by the release of these documents are rooted in political and PR considerations, and in EPA's manifest displacement of its statutory missions with promoting the interests of AI companies. Those harms are not cognizable under the deliberative process privilege.

Conclusion

Because chemical reviews affect public health, EPA should immediately provide answers about its AI programs and demonstrate that its governance framework is in fact real. This will allow independent scrutiny of the methodologies being used by the agency to review chemicals and help maintain trust in agency decisions.

Thank You,

Kevin H. Bell
Partner
Free Information Group, PLLC
kevin@freeinformationgroup.com

Counsel for Requester-Appellant, Public Employees for Environmental Responsibility

[16] H.R. REP. NO. 391, 114th Cong., 2d Sess. 10 (2016).