epic.org

Electronic Privacy Information Center
1519 New Hampshire Avenue NW
Washington, DC 20036, USA

+1 202 483 1140
+1 202 483 1248
@epic.org
https://epic.org

COMMENTS OF THE ELECTRONIC PRIVACY INFORMATION CENTER AND THE CENTER FOR DEMOCRACY & TECHNOLOGY

To the

DEPARTMENT OF HOUSING AND URBAN DEVELOPMENT

Notice of a Modified System of Records: Privacy Act of 1974, Customer Relationship Management System

91 Fed. Reg. 17,295 | FR–7106–N–15

May 6, 2026

The Electronic Privacy Information Center (EPIC) and the Center for Democracy & Technology (CDT) submit these comments in response to the Department of Housing and Urban Development (HUD or the Department)'s Notice of intent to modify the System of Records, "Customer Relationship Management" (CRM) (hereinafter the "Notice"), published on April 6, 2026.

EPIC is a public interest research center in Washington, D.C., established in 1994 to focus public attention on emerging civil liberties issues and to secure the fundamental right to privacy in the digital age for all people through advocacy, research, and litigation.[1] For decades, EPIC has engaged with federal agencies to safeguard individuals' privacy and human rights.[2] CDT is a nonprofit 501(c)(3) organization that focuses on advancing civil rights and civil liberties in the digital age. As AI use across federal agencies has escalated, EPIC and CDT have consistently called for transparent, equitable, commonsense AI policy and regulations to protect individuals against AI systems being used in harmful, opaque, and unaccountable ways.[3] Among the proposed changes are the collection of new categories of Personally Identifiable Information (PII) and an introduction of AI features through Microsoft Dynamics (hereinafter the "AI system") that "would use existing case/service

---

[1] EPIC, *About Us* (2025), https://epic.org/about/.
[2] *See, e.g.,* EPIC Comments to DHS on RFI on Accelerating the Adoption and Use of Artificial Intelligence as Part of Clinical Care (Feb. 23, 2026), https://epic.org/documents/epic-comments-to-hhs-re-rfi-ai-in-clinical-carea/; EPIC Comments to HUD on Implementation of the Fair Housing Act's Disparate Impact Standard (Feb. 13, 2026), https://epic.org/documents/epics-comment-to-hud-on-the-disparate-impact-rule/; Comments of EPIC to the White House OSTP, Request for Information: Automated Worker Surveillance and Management (June 15, 2023), https://epic.org/documents/epic-comments-to-white-house-ostp-on-automated-worker-surveillance/.
[3] *See, e.g., AI & Human Rights*, EPIC (2026), https://epic.org/issues/ai/; *AI Policy & Governance*, CDT (2026), https://cdt.org/area-of-focus/ai-policy-governance/; *AI Policy*, EPIC (2026), https://epic.org/issues/ai/ai-policy/; EPIC Urges the OSTP to Focus on AI Protections, Not Deregulation, EPIC (Oct. 28, 2025), https://epic.org/epic-urges-the-ostp-to-focus-on-ai-protections-not-deregulation/.

request information and the organization's internal knowledge base to generate draft case summaries and suggested email responses."[4]

While the Notice does not introduce new routine uses, the collection of new categories of data and use of AI will potentially impact every existing use of data in the system. Despite the potentially serious consequences of these proposed modifications, the Notice lacks significant details about system accuracy, fitness and safety, security protocols to protect PII, and accountability mechanisms for individuals whose PII is processed by the AI system. To make matters worse, the Notice suggests that these practices have already been taking place prior to publication of the Notice. These deficiencies violate the Privacy Act and thus EPIC and CDT urge the Department to roll back the changes proposed in the Notice and proceed with the proposed changes only after addressing the issues raised.

I. The Notice Lacks Sufficient Justification for the Collection of New Categories of Data

The Notice states that the Department will collect new categories of PII for two out of the three systems included in the SORN. These additional data categories include, for the Microsoft Dynamics CRM, "client type, client's company (if they are associated with HUD grantees or with external stakeholder entities), and service request categorization, which includes primary and sub-category, along with case modifiers such as disabled, elderly, or veteran."[5] For HUD Central, the Department will newly collect "fax number, user ID(s)," "home/work address(es), investigation report or database IP/MAC address, case number(s), property zip code(s), and speech-to-text."[6]

The Privacy Act imposes rigorous data minimization standards on federal agencies. An agency may only collect information for a purpose that either a statute or executive order require the agency to accomplish.[7] Even then, the agency may only collect information that is relevant and necessary to accomplish the identified purpose.[8] The Department completely fails to support why those new categories of PII are "relevant" or "necessary" to accomplish an identified purpose as required under the Privacy Act.

The identified purpose for the system is too vague to support an assertion that the collection of new data is "relevant" or necessary." The SORN states that "[t]his CRM System SORN combines the 'One Stop Customer Service, HUD Central, and Microsoft Dynamics' documenting the use of Customer Relationship Management systems to manage, track, route, and respond to interactions with all customers, stakeholders, partners, and organizations who initiate a customer service interaction with the Department."[9] The Notice also states that the additional PII are collected "to support the processing of inquiries from program participants and citizens."[10] The Notice provides no further explanation as to why the new categories of information, which include sensitive

---

[4] 91 Fed. Reg. 17295, 17295 (April 4, 2026).
[5] *Id*. at 17296.
[6] *Id*.
[7] 5 U.S.C. § 552a(e)(1).
[8] *Id*.
[9] 91 Fed. Reg. 17295, 17295 (April 4, 2026).
[10] *Id*.

information about individuals such as "disabled, elderly, or veteran," are necessary and relevant to newly collect, though the system presumably had functioned before without such information. New PII collection for a pre-existing system demands specific justification for why that new information is necessary to accomplish HUD's goals. Furthermore, the system already has 11 pre-existing routine uses[11] for which the new categories of data will be used. The Department is attempting to collect new PII that may be shared in various ways without a clear justification, flouting its responsibility for data minimization under the Privacy Act.

One category of newly collected information deserves heightened scrutiny—"speech-to-text," to be collected for HUD Central.[12] To define "speech-to-text" as a category of data is misleading, because it can capture any data included in the speech that is being processed into text, including sensitive data subject to higher protection requirements. In addition to the lack of justification for collecting what could be a large amount of data, the Notice provides no information as to the source of the data, the tools used to transform speech into text, and to what end the information will be used. Would individuals who call a HUD office for any reason have their speech be transcribed into text and recorded? Would those individuals be informed of this and have an option to refuse and still receive services? An individual may divulge information in this context that goes beyond the identified data categories that the system collects. The failure to provide justification for its collection and lack of clarification on the scope of speech-to-text data collection make this Notice clearly deficient under the Privacy Act.

II. The Notice Lacks Assurance of Accuracy and Fairness for the Modified System, Violating the Privacy Act

The Privacy Act contains several provisions designed to ensure that the records collected about individuals are accurate, relevant, timely, and complete.[13] The collection of speech-to-text data and the use of the AI system raise concerns that the Department is falling short of the Privacy Act's requirements for accuracy and fitness.

Research shows that speech-to-text systems can produce errors and biases. Specifically, studies of speech recognition technologies show that the systems have higher error rates for racial minorities, even when the speakers were matched by age and gender and spoke the same words.[14] The studies also show that speech recognition systems perform worse on speakers whose voice characteristics deviate from the norm, such as Deaf and Hard of Hearing speakers, and show biases around racial,

---

[11] *Id*. at 17296–97.
[12] 91 Fed. Reg. 17295, 17296 (April 4, 2026).
[13] 5 U.S.C. § 552a(e)(5) (requiring agencies to "maintain all records which are used by the agency in making any determination about any individual with such accuracy, relevance, timeliness, and completeness as is reasonably necessary to assure fairness to the individual in the determination."); 5 U.S.C. § 5a(e)(6) (requiring that "prior to disseminating any record about an individual to any person other than an agency, unless the dissemination is made pursuant to subsection (b)(2) of this section, make reasonable efforts to assure that such records are accurate, complete, timely, and relevant for agency purposes.").
[14] Edmund L. Andrews, *Stanford researchers find that automated speech recognition is more likely to misinterpret black speakers*, Stanford Report (March 23, 2020), https://news.stanford.edu/stories/2020/03/automated-speech-recognition-less-accurate-blacks#:~:text=Mis%2Dtranscribed%20speech&text=All%20five%20speech%20recognition%20technologies,of%20samples%20spoken%20by%20whites.

ethnic, and dialect distinctions, including misrecognizing speech from speakers of African American Vernacular English (AAVE).[15] Against this backdrop, the lack of clarity from the Department on the accuracy of the speech-to-text data that it intends to collect brings into question whether it is complying with the Privacy Act's requirement for ensuring accuracy of the records.

The proposed AI system use casts doubt on whether the Department would comply with the Privacy Act's requirement for accuracy of records and fairness when making determinations about an individual.[16] The SORN states that the AI system will "use existing case/service request information and the organization's internal knowledge base to generate draft case summaries and suggested email responses."[17] The Notice further states "[a]ll record categories will support resolving a contact's issue by enabling AI to summarize cases, assign priority levels, and generate draft email responses from internal knowledgebase content, with all AI-produced drafts reviewed by staff before sending."[18] This raises several red flags.

First, generative AI systems like the one that the Department is proposing to use produce hallucinations that would undermine the accuracy of records.[19] Second, using AI systems in making decisions that require human judgment, such as those used to produce automated priority levels and draft decisions, come with risks of biases, inaccuracies, and unfair decision making. Lastly, the introduction of an AI system may increase, not decrease, burden on staff and individuals whose records the Department collects and processes.

With the AI system, the Department introduces a new way to create inaccurate records through case summaries or draft emails. Put simply, generative AI tools are built to calculate the next word that is most statistically plausible to produce outputs given the data on which the model was trained.[20] This means that the system itself does not understand the meaning of the words it is stringing together or whether the output is true or false. As a result, generative AI systems inevitably produce seemingly plausible but inaccurate outputs, or what are often referred to as hallucinations.[21] The Notice fails to address the possibility of hallucinations, any audits to be performed on the accuracy of the AI system (both prior to use and throughout use), or how it will minimize the risk of introducing inaccurate information through the AI system, as is required to comply with the Privacy Act.

---

[15] Korbinian Kuhn, Verena Kersken, Benedikt Reuter, Niklas Egger & Gottfried Zimmermann, *Measuring the Accuracy of Automatic Speech Recognition Solutions*, arXiv (Aug. 2024), https://arxiv.org/abs/2408.16287; Mikel K. Ngueajio & Gloria Washington, *Hey ASR System! Why Aren't You More Inclusive? Automatic Speech Recognition Systems' Bias and Proposed Bias Mitigation Techniques*. *A Literature Review.,* arXiv (Nov. 2022), https://arxiv.org/html/2508.07143v1#:~:text=A%20growing%20body%20of%20research,(Zhao%20et%20al.%20)%20.

[16] 5 U.S.C. § 552a(e)(5).

[17] 91 Fed. Reg. 17295, 17296 (April 4, 2026).

[18] *Id*.

[19] Conor Murray, *Why AI 'Hallucinations' Are Worse Than Ever,* Forbes (May 6, 2025), https://www.forbes.com/sites/conormurray/2025/05/06/why-ai-hallucinations-are-worse-than-ever/.

[20] Cole Stryker, *What are large language models (LLMs)?*, IBM, https://www.ibm.com/think/topics/large-language-models; Emily M. Bender, Timnit Gebru, Angelina McMillan-Major & Shmargaret Shmitchell, *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610–23 (March 1, 2021), https://doi.org/10.1145/3442188.3445922.

[21] Gyana Swain, *OpenAI admits AI hallucinations are mathematically inevitable, not just engineering flaws*, ComputerWorld (Sept. 18, 2025), https://www.computerworld.com/article/4059383/openai-admits-ai-hallucinations-are-mathematically-inevitable-not-just-engineering-flaws.html; Ziwei Xu, Sanjay Jain & Mohan Kankanhalli, *Hallucination is Inevitable: An Innate Limitation of Large Language Models*, arXiv (Feb. 13, 2025), https://arxiv.org/abs/2401.11817.

The Department's insistence that "all AI-produced drafts [are] reviewed by staff before sending" and that the AI system "does not independently make determinations about individuals" do little to assure the public on accuracy of maintained records and fairness of determinations.[22] Not only do generative AI systems produce inaccurate outputs, they do so confidently.[23] Studies show that humans often over-rely on AI system outputs, making human review not a reliable fix for erroneous AI outputs.[24] There are numerous real-world examples that show that introduction of AI systems to automate decisionmaking, even with human review, perpetuate erroneous and biased outputs, causing immense harm. Improper automated decisionmaking have already become a notable problem in both housing benefits[25] and other critical areas, such as health care,[26] job opportunities,[27] public benefits,[28] and more.[29] Similar to the Notice, companies using AI systems to process health insurance claims insist that a human reviews AI system recommendations and that humans make the final decision.[30] However, data shows that human reviewers of medical claims only "review" AI recommendations for mere seconds before approving them, a seemingly impossible time frame to validate the outputs.[31] In cases like this, the human review is merely a rubber stamp. The Notice fails to clarify how HUD will ensure that human review is meaningful and will account for AI errors, misinterpreting data, context, and using their own expertise and judgment.

While the Department claims that the AI system will be used to reduce administrative burdens, the AI system may instead increase the burdens on staff and individuals who engage with HUD. Various studies show that companies have not seen return on investment on the deployment of AI, with many workers complaining of AI slop, which are AI-produced work products that seem credible on first skim but lack accuracy and substance.[32] Some workers also report that their expertise is undermined

---

[22] 91 Fed. Reg. 17295, 17296 (April 4, 2026).
[23] Celina Zhao, *AI Hallucinates Because it's Trained to Fake Answers it Doesn't Know*, Science (Oct. 28, 2025), https://www.science.org/content/article/ai-hallucinates-because-it-s-trained-fake-answers-it-doesn-t-know.
[24] Neil Rathi, Dan Jurafsky & Kaitlyn Zhou, *Humans Overrely on Overconfident Language Models, Across Languages*, arXiv (July 8, 2025), https://arxiv.org/html/2507.06306v1; Eric Bogert, Aaron Schecter & Richard T. Watson, *Humans Rely More on Algorithms Than Social Influence as a Task Becomes More Difficult*, 11 Scientific Reports 8028 (2021), https://www.nature.com/articles/s41598-021-87480-9.
[25] Johana Bhuiyan, *She Didn't Get an Apartment Because of an AI-Generated Score – and Sued to Help Others Avoid the Same Fate*, Guardian (Dec. 14, 2024), https://www.theguardian.com/technology/2024/dec/14/saferent-ai-tenant-screening-lawsuit.
[26] Ziad Obermeyer, Brian Powers, Christine Vogeli & Sendhil Mullainathan, *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, Science (Oct. 25, 2019), https://www.science.org/doi/10.1126/science.aax2342.
[27] Charlotte Lytton, *AI Hiring Tools May Be Filtering out the Best Job Applicants*, BBC (Feb. 16, 2024), https://www.bbc.com/worklife/article/20240214-ai-recruiting-hiring-software-bias-discrimination.
[28] Rachana Pradhan, Samantha Liss & KFF Health News, *A Tennessee Mom Lost Medicaid After the State Launched a Deloitte-Run System that Managed Eligibility. Then Her Life Turned Upside Down*, Fortune (June 24, 2024), https://fortune.com/2024/06/24/a-tennessee-mom-lost-medicaid-after-the-state-launched-a-deloitte-run-system-that-managed-eligibility-then-her-life-turned-upside-down/.
[29] Mayu Tobin-Miyaji, EPIC, *Assessing the Assessments: Maximizing the Effectiveness of Algorithmic & Privacy Risk Assessments* 5–15 (2025),
[30] *See* Patrick Rucker, Maya Miller, and David Armstrong, *How Cigna Saves Millions by Having Its Doctors Reject Claims Without Reading Them,* ProPublica (Mar. 25, 2023), https://www.propublica.org/article/cigna-pxdx-medical-health-insurance-rejection-claims#:~:text=But%20the%20Cigna,at%20a%20time.%E2%80%9D.
[31] *Id*.
[32] Jason Koebler, *AI 'Workslop' Is Killing Productivity and Making Workers Miserable*, 404 Media (Sept. 23, 2025), https://www.404media.co/ai-workslop-is-killing-productivity-and-making-workers-miserable//.

when employers require AI because they are directed to follow AI outputs that may be inaccurate and overriding the system's outputs is difficult or discouraged.[33] HUD staff may find themselves with more work, not less, in constantly reviewing and correcting outputs from the AI system. In addition, the Privacy Act gives individuals the right to access records held about them by the Department and request to correct records if they are inaccurate.[34] AI systems' likelihood of producing inaccurate outputs may increase the burden on the members of the public that interact with HUD by forcing them to spend time and effort to regularly exercise those rights in correcting inaccurate data.

III. The Deployment of the AI System is Not Compatible with the Routine Uses Under the Privacy Act

The Privacy Act generally prohibits disclosure of personal records unless exceptions apply. Among the listed exceptions to the disclosure prohibition are "routine use" disclosures, i.e., those made "for a purpose which is compatible with the purpose for which [a record] was collected."[35] Although all routine uses listed in the Notice are the same as those listed in the previous SORN,[36] the introduction of the AI system changes the nature of the routine uses and no longer satisfies the compatibility requirement under the Privacy Act.

Determining "compatibility" under the Privacy Act requires "a dual inquiry into the purpose for the collection of the record in the specific case and the purpose of the disclosure."[37] In order to qualify for the exception, a mere "relevance" of the disclosure's objective to the collection's aim is insufficient—"[t]here must be a more concrete relationship or similarity, some meaningful degree of convergence, between the disclosing agency's purpose in gathering the information and in its disclosure."[38] In other words, the purpose of information collection and the purpose of the disclosure must be strongly connected.

The purpose of the system, according to the Notice, is "to manage, track, route, and respond to interactions the Department has with the public, stakeholders, partners, and other organizations interested in how HUD does business, such as advocacy groups, professional organizations, Congress, and the media."[39] The Notice explains that "AI features would use existing case/service request information and the organization's internal knowledge base to generate draft case summaries and suggested email responses. All record categories will support resolving a contact's issue by enabling AI to summarize cases, assign priority levels, and generate draft email responses from internal knowledgebase content, with all AI-produced drafts reviewed by staff before sending. It is designed to reduce administrative burden, improve consistency in communications, and help staff quickly identify relevant information when resolving customer issues."[40]

---

[33] Press Release, *National Nurses United survey finds A.I. technology degrades and undermines patient safety*, National Nurses United (May 15, 2024), https://www.nationalnursesunited.org/press/national-nurses-united-survey-finds-ai-technology-undermines-patient-safety.
[34] 5 U.S.C. § 552a(d)(1)–(2).
[35] 5 U.S.C. §§ 552a(a)(7), (b)(3).
[36] 89 Fed. Reg. 86352 (Oct. 30, 2024).
[37] *Britt v. Naval Investigative Serv.*, 886 F.2d 544, 548–49 (3d Cir. 1989).
[38] *Id*. at 549–50 (citing Mazaleski v. Treusdell, 562 F.2d 701, 713 n. 31 (D.C.Cir. 1977)) (other citation omitted).
[39] 91 Fed. Reg. 17295, 17296 (April 4, 2026).
[40] *Id*.

The deployment of the AI system to ingest "all record categories" to enable its AI uses is incompatible with the purpose for the collection of the records. Even compared against the vague and insufficient justification of the collection of additional PII "to support the processing of inquiries from program participants and citizens," the use of the AI system in any of the routine purposes is incompatible. The Department's deployment of the AI system will lead to the AI system using all data points collected by the CRM system or held by the Department as training data to generate outputs for any of the routine uses. Training an AI system is a completely distinct use from the purpose for which the data was originally collected, failing the compatibility requirement. The Notice merely cites to efficiency and burden justifications, but those considerations do not create the level of "meaningful degree of convergence" to ensure compatibility of purposes under the Privacy Act. The Department is sidestepping its responsibility under the Privacy Act.

The bias and inaccuracy concerns raised by AI systems and the human tendency to over-rely on AI outputs also make its deployment incompatible under the Privacy Act. As discussed above, there are real concerns that the AI system will generate biased or inaccurate outputs that are not sufficiently mitigated by human reviewers and impact determinations. The Department takes the position that "all AI-produced drafts are reviewed by staff before sending" and that the AI system "does not independently make determinations about individuals." However, existing research shows that the bias and accuracy issues are difficult to resolve through human oversight. The Notice also does not specify whether the AI system will be audited for accuracy or bias issues, whether the AI-generated and human-reviewed outputs will be recorded separately, and whether divergence in these outputs will be reviewed to assess impact on determinations. Thus, the use of the AI system to process all collected data for every routine use, including to aid in making determinations, fails to meet the compatibility standard under the Privacy Act.

IV. Several Listed Routine Uses Become Higher-Risk When Combined with AI Systems

The Notice does not identify whether or how information used for Routine Use 9, to "assist in the enforcement of civil or criminal laws,"[41] will be reviewed for accuracy, including which other records the information may be combined with. When combined with other sources of information, the records covered by the Notice may yield inaccuracies or errors that could lead to the deprivation of rights in a law enforcement setting. This may be due to mismatched records, minor errors in records that become consequential given heightened risk of law enforcement use, or confusion stemming from using information outside its original context. The Notice's failure to identify a process for ensuring the accuracy and completeness of records thus violates the Privacy Act and threatens to undermine the fairness and credibility of law enforcement activities.

The Notice includes an overly broad routine use, Routine Use 6, that would permit the Department to share information with "contractors, experts, and consultants…for the purpose of testing new technology and systems designed to enhance program operations and performance."[42] This language seems to allow for any AI model training with the data. The use of the AI system in conjunction with this routine use is incompatible with the original purpose of collection as it is well outside of the

[41] *Id*. at 17297.
[42] *Id*. at 17296.

system's purpose to "manage, track, route, and respond to interactions the Department has with the public, stakeholders, partners, and other organizations."[43]

Lastly, the Notice does not state whether Microsoft will be able to use this data to enrich their models generally or whether the data and algorithms built off of that PII will be wholly siloed within HUD. Allowing Microsoft to profit from this PII by enriching their broader models not only increases security risk to individuals with implicated data, but it also exploits the PII far beyond the explicit purposes listed within HUD policies and the Notice itself. Individuals provide this data to receive services from HUD, and allowing Microsoft to repurpose that data to train its own AI models would be completely incompatible with the original purpose of the data collection and a clear violation of the Privacy Act.

V. Cybersecurity and Data Protection Requirements Under the Privacy Act are Left Unaddressed

The Privacy Act requires agencies to establish safeguards to ensure the security and confidentiality of records.[44] The use of the AI system heightens data security risks and the Notice fails to provide any mention of necessary safeguards. The Notice states that all data categories are used in the AI system, meaning that the AI system trains on all PII collected. PII fed into AI systems may be revealed if the system is breached, as has occurred before.[45] A system trained on all data from HUD's communications with external entities creates a rich target for unauthorized access, heightening the risk of a breach. HUD fails to address such concerns.

The nature of AI systems to retain insights from all training data also contradicts the Notice's retention and disposal practices. The Notice outlines the disposal schedule of data after a set amount of time. However, an AI model can retain insights from training data indefinitely, even if the data itself is deleted, and expose the sensitive training data embedded into the system.[46] Complete deletion of the PII would require retraining the AI model without the PII in the training dataset and the Notice does not indicate that the Department plans to do so.[47] The Notice as it stands misleads the public on its data retention practices, again violating the Privacy Act.

---

[43] *Id*.
[44] 5 USC 552a(e)(10).
[45] *See* Milad Nasr, Nicholas Carlini, Katherine Lee, et. al, *Extracting Training Data from ChatGPT,* arXiv (Nov. 28, 2023); Comments of EPIC to the White House OSTP, Request for Information: Automated Worker Surveillance and Management (June 15, 2023), https://not-just-memorization.github.io/extracting-training-data-from-chatgpt.html; Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, Alina Oprea & Colin Raffel, *Extracting Training Data from Large Language Models,* arXiv (Dec. 4, 2020), https://arxiv.org/abs/2012.07805#:~:text=in%20such%20settings%2C%20an%20adversary,document%20in%20the%20training%20data.
[46] Ruth Ntumba, *AI models may retain and expose sensitive personal data despite industry safeguards, study finds*, Imperial College London (March 19, 2026), https://www.imperial.ac.uk/news/articles/engineering/computing/2026/ai-has-a-memory-problem--and-your-work-data-may-be-at-risk/.
[47] Antonio A. Ginart, Melody Y. Guan, Gregory Valiant & James Zou, *Making AI Forget You: Data Deletion in Machine Learning*, 33rd Conference on Neural Information Processing Systems (2019), https://proceedings.neurips.cc/paper_files/paper/2019/file/cb79f8fa58b91d3af6c9c991f63962d3-Paper.pdf.

VI. The Notice Violates the Privacy Act's Requirement to Publish Notice Before Changes are Made

Lastly, we object to the Department's apparent after-the-fact notice of modifications to the SORN. Belatedly providing notice of changes already made flouts the Privacy Act's requirement to notify the public and conduct a meaningful review before implementing changes.

The Privacy Act requires an agency to, "at least 30 days prior to publication of information under paragraph (4)(D) of this subsection, publish in the Federal Register notice of any new use or intended use of the information in the system, and provided an opportunity for interested persons to submit written data, views, or arguments to the agency."[48] Paragraph (4)(D) of the subsection refers to "each routine use of the records contained in the system, including the categories of users and purposes of such use."[49] Further, agencies cannot "use a new or significantly modified routine use as the basis for a disclosure fewer than 30 days following Federal Register publication."[50]

Several passages in the Notice, however, suggest that the Department did not comply with the pre-deployment notice requirement under the Privacy Act. For example, the Notice states, "[t]he reason for this revised notice is to make clarifying changes within: System Manager, and Categories of Records to update description of data collection activities performed by the One Stop Customer Service, HUD Central, and Microsoft Dynamics systems, and Policies and Practices for Retention and Disposal for Records to add context."[51] That the Notice is "clarifying changes…to add context" suggests that the changes have already taken place, and the Department is publishing this Notice to merely bring the paperwork up to date. Further, the Notice states that "the 'Categories of Records in the system' section has been updated to include record categories that were previously omitted from the published SORN."[52] There should be no such thing as record categories that were "previously omitted" from a SORN that were nonetheless collected by the Department. That would constitute a clear violation of the Privacy Act. Furthermore, the Notice includes, "[t]his modification also updates contact information in the 'System Manager(s)' section to reflect recent personnel changes and revises the 'Policies and Practices for Retention and Disposal for Records' to provide additional context and ensure the section reflects the most current information."[53] Once again, "to provide additional context and ensure the section reflects the most current information" strongly suggests that these changes have already been made, and this Notice constitutes an update to the SORN after the fact. The Privacy Act forbids such ex-post notification to the public.

VII. Conclusion

For the above reasons, HUD should promptly withdraw its Notice and roll back any changes made before the Notice was published in violation of the Privacy Act. HUD should proceed with the proposed changes only after addressing the issues raised and take the steps necessary to ensure that

---

[48] 5 U.S.C. § 552a(e)(11).
[49] *Id*. (e)(4)(D).
[50] Off. of Mgmt. & Budget Circular No. A-108, Federal Agency Responsibilities for Review, Reporting, and Publication under the Privacy Act, at 7 (2016), https://perma.cc/N9QK-SDLE.
[51] 91 Fed. Reg. 17295, 17295 (April 4, 2026).
[52] *Id*.
[53] *Id*. at 17295–96.

any changes to be made to the CRM system comply with the Privacy Act's requirements. It should go without saying that if the proposed collection of new categories of data and use of the AI system cannot comply with the Privacy Act's requirements, the Department must not pursue that path.

Respectfully submitted,

/s/ Calli Schroeder
Calli Schroeder
Senior Counsel and Director, AI and Human Rights Program
ELECTRONIC PRIVACY INFORMATION CENTER (EPIC)

/s/ Mayu Tobin-Miyaji
Mayu Tobin-Miyaji
EPIC Law Fellow
ELECTRONIC PRIVACY INFORMATION CENTER (EPIC)

/s/ Elizabeth Laird
Elizabeth Laird
Director of Equity in Civic Technology
CENTER FOR DEMOCRACY & TECHNOLOGY (CDT)

/s/ Quinn Anex-Ries
Quinn Anex-Ries
Senior Policy Analyst
CENTER FOR DEMOCRACY & TECHNOLOGY(CDT)