



April 29, 2024

To: Samantha Hubner

Office of the Federal Chief Information Officer, Office of Management and Budget
725 17th Street NW
Washington, DC 20503

Submitted via regulations.gov

Re: Request for Information on Responsible Procurement of Artificial Intelligence in Government, OMB-2024-0004

The Center for Democracy & Technology (CDT) respectfully submits these comments in response to the Office of Management and Budget (OMB) Request for Information on responsible procurement of artificial intelligence (AI) in government. CDT is a nonprofit 501(c)(3) organization that works to advance civil rights and civil liberties in the digital age. CDT's work includes advocating for responsible government procurement and use of technology, including AI, to ensure that government services and benefits are delivered efficiently while protecting people's privacy and civil rights.

CDT's new report, *The Federal Government's Power of the Purse: Enacting Procurement Policies and Practices to Support Responsible AI Use*,¹ which are attached to these comments, addresses several of the questions set forth in OMB's request for information. Therefore, our comments include relevant excerpts from this report that identify recommendations for OMB to incorporate into its "initial means to ensure that agency contracts for the acquisition of AI systems and services align with" OMB M-24-10. In addition to the recommendations in our own report, we encourage OMB to refer to the comments by EPIC, the Brennan Center, et al. on this Request for Information, as they suggest additional avenues for building an effective framework for procuring responsible and rights-respecting AI systems.² Similarly, we encourage OMB to refer to comments by the Lawyers' Committee on Civil Rights Under Law regarding

¹ See Center for Democracy & Technology, *The Federal Government's Power of the Purse: Enacting Procurement Policies and Practices to Support Responsible AI Use* (2024), <https://cdt.org/insights/report-the-federal-governments-power-of-the-purse-enacting-procurement-policies-and-practices-to-support-responsible-ai-use/> [hereinafter "The Federal Government's Power of the Purse"].

² See Electronic Privacy Information Center et al, *Comment to Office of Management and Budget on Responsible Procurement of Artificial Intelligence in Government* (Apr. 29, 2024).

advancing equitable outcomes via AI procurement, which highlight the pressing need for improved AI procurement frameworks for the federal government to advance equitable outcomes and protect civil rights.³

1. How may standard practices and strategies of federal procurement be best used to reflect emerging practices in AI procurement? Are there additional materials or resources that OMB could provide to vendors or agencies to improve alignment and between agency missions and technical requirements?

Part V of *The Federal Government's Power of the Purse* makes a number of recommendations for how the federal government can use or build on existing standard practices or requirements in federal procurement. Pertinent highlights include the following:

Incorporate an “AI Responsibility Questionnaire” into the Requests for Information that agencies often direct to vendors for market research:⁴ “When conducting market research as part of acquisition planning...[a]n AI Responsibility Questionnaire could help formalize some of the considerations that agencies should take into account, and could be developed in modular fashion to suit different agencies’ fact patterns and contexts. A questionnaire can also establish expectations for industry about the types of questions that companies who plan to sell to the government should be able to answer.” Ensuring that agencies actually benefit from this tool “will require advocacy and education for public administrators to be able to use the questionnaire effectively, and make better decisions based upon it.”

Require legal review of all contracts involving AI.⁵ “Contracting offices have standard operating procedures concerning internal approvals to address risks. Often, internal approvals are threshold-based (e.g., contracts in excess of \$1 million) or trigger-based (e.g., sole-source contracts).

Given the novelty of legal risks and harms associated with AI, including copyright issues, privacy concerns, and the unique impact that vendors’ standard intellectual property clauses may have on agencies’ ability to explain and test AI systems, contracts involving AI should be reviewed by agency counsel with relevant expertise.”

³ See Lawyers’ Committee for Civil Rights Under Law, Comment to Office of Management and Budget on Responsible Procurement of Artificial Intelligence in Government (Apr. 29, 2024).

⁴ See *The Federal Government's Power of the Purse* at 33 (citations omitted).

⁵ See *The Federal Government's Power of the Purse* at 34 (citations omitted).

Explicitly call out responsible AI practices in the Federal Acquisition Regulations (FAR).⁶ “Any amendments [to the FAR] would likely be useful if they provide specific references back to requirements in [Executive Order 14110, OMB Memorandum M-24-10, Executive Order 14091], and any applicable statutes, as well as guidance provided in NIST’s AI Risk Management Framework and the Blueprint for an AI Bill of Rights. Additionally, because of the specific risks associated with ‘drift’ (the ongoing efficacy and suitability of an AI model over time), there may be particularly useful changes in Part 46 (Quality Assurance) because traditional rules around acceptance – ‘acknowledgement that the supplies or services conform with applicable contract quality and quantity requirements’ – may not be sufficient to manage AI risks.”

Add specific reporting requirements on responsible AI in the Federal IT Acquisition Reform Act (FITARA) scorecard.⁷ “FITARA has been a major driver in moving agencies to adopt the commercial cloud because of the visibility of the FITARA scorecard, which uses a simple grading method to track how well agencies are complying with FITARA’s requirements to manage their IT adoption. Including metrics that reflect responsible AI practices in the FITARA scorecard (e.g., ‘percent of systems that have independent algorithmic impact assessments’) could induce CIOs to prioritize responsible AI practices and policies... Further, [Executive Order 14110] requires OMB to develop a method for agencies to ‘track and assess their ability to adopt AI into their programs and operations [and] manage its risks,’ addressing the practices and capabilities needed for AI governance across IT infrastructure, risk management, and other areas.

Develop oversight guidance that builds on GAO’s *Accountability Framework for Federal Agencies and Other Entities*, to identify questions for federal employees and contractors.⁸ “Building on the GAO’s *Accountability Framework for Federal Agencies and Other Entities*, the Congressional Research Service or agencies responsible for oversight like GAO or OMB should develop an ‘oversight guide’ with specific references to parts of the acquisition planning process where agencies should *already* be considering responsible AI principles.

An oversight guide could be used by Congress, federal agency executives, inspectors general, civil society, and the public to better understand what types of questions to ask of federal government employees and federal contractors. Because structured oversight processes can help drive agencies toward better outcomes, having a shared understanding of what ‘good’ looks like from an oversight perspective would be useful.”

⁶ See *The Federal Government’s Power of the Purse* at 35 (citations omitted).

⁷ See *The Federal Government’s Power of the Purse* at 43-44 (citations omitted).

⁸ See *The Federal Government’s Power of the Purse* at 45 (citations omitted).

Perform AI impact assessments in addition to privacy impact assessments during the procurement process.⁹ “Agencies are required by law to conduct Privacy Impact Assessments (PIAs) when ‘developing or procuring information technology that collects, maintains, or disseminates information that is in an identifiable form’ and publish them on the agency’s website...OMB has advised that a PIA must be written in ‘plain language’ and address ‘how information is handled: (i) to ensure handling conforms to applicable legal, regulatory, and policy requirements regarding privacy, (ii) to determine the risks and effects of collecting, maintaining and disseminating information in identifiable form in an electronic information system, and (iii) to examine and evaluate protections and alternative processes for handling information to mitigate potential privacy risks.’”

Similarly, agencies should conduct algorithmic impact assessments (AIAs)—either as part of the PIA process or as a standalone assessment—whenever an automated system is used to make or influence a decision that has a service-delivery impact. [OMB Memorandum M-24-10] requires AI impact assessments as a minimum risk management practice for rights- or safety-impacting systems (covering the majority of instances where a PIA would be required).”

Work with Digital IT Acquisition Profession (DITAP) program providers to incorporate responsible AI concepts into the DITAP curriculum.¹⁰ “Working with United States Digital Service (USDS), the Office of Federal Procurement Policy (OFPP), and DITAP providers to incorporate responsible AI concepts into the DITAP curriculum would have a broad reach among acquisition professionals who plan to work with digital services teams.

Similarly, civil society organizations also offer training programs that could benefit from the inclusion of information about how federal leaders can responsibly procure AI. Partnering with these organizations to incorporate responsible AI training into the Federal Acquisition Institute curriculum will help reach acquisition professionals (including Contracting Officers and their representatives).”

2. How can OMB promote robust competition, attract new entrants (including small businesses) into the federal marketplace, and avoid vendor lock-in across specific elements of the technology sector, including data collectors and labelers, model developers, infrastructure

⁹ See *The Federal Government’s Power of the Purse* at 46-47 (citations omitted).

¹⁰ See *The Federal Government’s Power of the Purse* at 49 (citations omitted).

providers, and AI service providers? Are there ways OMB can address practices that limit competition, such as inappropriate tying, egress fees, and self-preferencing?

OMB can best and most appropriately promote competition in federal AI procurement through the federal government's own direct action as the consumer.

OMB should develop an AI procurement policy that makes creation of opportunities for new entry and competition an explicit consideration in each agency's AI procurement process. That should begin with careful specification of the procurement requirements so as not to inadvertently exclude goods and services that may be effective substitutes. The procurement process should also not include criteria that may not be necessary but that could have the effect of raising entry barriers. For example, a procurement specification that called for an AI model of a certain size or compute power may have the effect of excluding smaller entrants even if a model of lower size or power could still provide the desired functionality.

OMB and the agencies should actively promote supplier entry to increase competition. It can provide guidance and support and resources (e.g., data sets that can be used for training).

The procurement process should stipulate that among the factors that the agency will weigh favorably in deciding among applicants for a procurement contract is an applicant's commitment to include independent partners and subcontractors to fulfill it. An applicant should be required to list all partners and subcontractors, any relationship they have to or history they have with the applicant, their contributions and proposed remunerations, their workforce size, their time in existence, and other attributes that OMB determines would help illuminate the breadth and inclusiveness of the participants the applicant would bring with it. Some of this information would be competitively sensitive and should be submitted in confidence, and not shared outside the agency officials involved in the procurement, subject to appropriate oversight in the agency and in the Antitrust Division.

Implementing this procurement policy successfully will require that OMB, and each agency, overcome the tendency to award the contract to the lowest qualified bidder as a matter of course. Or stated another way, require that OMB make promotion of competition a strong enough consideration to overcome cost as the predominant determinant in appropriate circumstances.

OMB and the agencies will also need to devote sufficient resources to ensure independent government assessment of the prospective partners and subcontractors. Reliability and safety cannot be compromised for the sake of wider inclusion.

And in developing this policy, OMB should actively seek out early and meaningful input from interested start-ups and small businesses. It can begin by following up with start-ups and small businesses that submit responses to this RFI. But it should go beyond that. The start-ups and small businesses OMB needs to hear from may not have the wherewithal to develop a timely response to the RFI, or even be aware of it. The leading AI developers will have no trouble getting their views considered. OMB will have to make an extra, proactive effort to ensure that the other views are also heard and fully considered.

The agencies can obtain further guidance generally from OIRA's October 2023 *Promoting Competitive Markets Through Regulatory Design*.¹¹ And when preparing specific contemplated procurement rulemakings, they can obtain guidance from the Justice Department's Antitrust Division under the interagency review and clearance process.

4. How might metrics be developed and communicated to enable performance-based procurement of AI? What questions should agencies be asking vendors to determine whether AI is already being used in performance-based services contracts?

As mentioned above, Part V of *The Federal Government's Power of the Purse* explains that an AI Responsibility Questionnaire could communicate expectations to industry about metrics that reflect the government's goals for the systems being procured, and the report recommends including metrics in the FITARA scorecard that reflect agencies' adoption of responsible AI practices.¹² In addition, "[a]lthough the FITARA scorecard is maintained by the House's Subcommittee on Government Operations, Committee on Oversight and Reform, the 'data used for grading federal agencies have largely been publicly available and regularly updated.' Agreeing upon what metrics should go into a scorecard would likely require coordination between the Subcommittee, GAO, and OMB. This may be a promising opportunity given GAO's

¹¹ Office of Information and Regulatory Affairs, *Promoting Competitive Markets Through Regulatory Design* (Oct. 11, 2023), <https://www.whitehouse.gov/omb/briefing-room/2023/10/11/promoting-competitive-markets-through-regulatory-design/>.

¹² See *The Federal Government's Power of the Purse* at 43-44 (citations omitted).

prior work on the AI Accountability Framework, and OMB’s current work on guidance for both agency use and agency procurement of AI.”

5. What access to documentation, data, code, models, software, and other technical components might vendors provide to agencies to demonstrate compliance with the requirements established in the AI M-Memo? What contract language would best effectuate this access, and is this best envisioned as a standard clause, or requirements-specific elements in a statement of work?

Vendors should be required to provide to agencies a number of technical components to demonstrate compliance with the risk management practices in the Final OMB AI Memo. Key among these components will be test results when the model is evaluated for performance, where vendors should provide at least:

- **Conceptual definitions and technical implementation details of evaluations:** Vendors should provide clear, conceptual definitions of the system capabilities and risks they have evaluated and how they have operationalized these concepts within their evaluation approach. For example, suppose a vendor offers an AI service that generates text and provides agencies with an evaluation that demonstrates the privacy risks of the system. This evaluation should be accompanied by a conceptual definition of what privacy means in this context, such as the extent to which the model might regurgitate sensitive content from its training data. The vendor should also provide a technical description of how this is operationalized in their evaluation, such as how often the model reproduces the rest of a text passage known to be in the training data when it is prompted with the first 50 tokens of that passage. Where vendors have evaluated AI system capabilities and risks using offline datasets (e.g., using accuracy metrics calculated on test data or using benchmark datasets), vendors should justify why these metrics are expected to generalize to the application deployed in context since offline performance and performance in deployment are not always correlated.¹³
- **Result summaries disaggregated per relevant factors:** Vendors should be able to disaggregate testing results by factors that may impact users’ rights, such as

¹³ See generally Joeran Beel and Stefan Langer, *A Comparison of Offline Evaluations, Online Evaluations, and User Studies in the Context of Research-Paper Recommender Systems*, in Proceedings of the 19th International Conference on Theory and Practice of Digital Libraries (2015), [https://docear.org/papers/A%20Comparison%20of%20Offline%20Evaluations,%20Online%20Evaluations,%20and%20User%20Studies%20...%20\(preprint\).pdf](https://docear.org/papers/A%20Comparison%20of%20Offline%20Evaluations,%20Online%20Evaluations,%20and%20User%20Studies%20...%20(preprint).pdf); Seraphina Goldfarb-Tarrant et al, *Intrinsic Bias Metrics Do Not Correlate with Application Bias*, arXiv (2021), <https://arxiv.org/pdf/2012.15859.pdf>.

demographic factors like race, gender, and disability status. Vendors should provide at least enough information for agencies to be confident that the system they are procuring performs comparably across demographic groups. In some cases, vendors may not have access to the demographic data they would need to perform these tests, but in this context vendors should still be able to provide information about steps they have taken to ensure equitable performance in the absence of such data. Additionally, agencies should include contract terms that require vendors to do demographic analyses during deployment and address any disparities that arise. Example language may be: “The vendor shall provide the results of any bias audits and demographic testing they have conducted on their product, as well as information about how any concerns or problems uncovered by the audit or testing have been addressed. If the vendor has not performed such an audit, they must do so within *[insert time period]* of receiving a contract with *[insert agency]* and will take steps to mitigate any issues uncovered within *[insert time period]* of completion of the audit.”¹⁴

- **Instance-by-instance evaluations:** Because even disaggregated, group-level metrics may not provide sufficient insight into how models will perform on relevant cases,¹⁵ vendors should provide agencies with greater granularity in their evaluation. If a model is evaluated on particular test sets or benchmarks, vendors should provide access to the evaluation set inputs, model outputs, and scoring for each individual instance, in addition to aggregate and disaggregated statistics of the model’s performance. This can help the agencies determine whether the aggregate statistics are even relevant. In many cases, certain specific results will be more relevant than the aggregate performance as some errors may be more impactful in a government use context than others. For example, some benchmarks evaluate language model bias by comparing the probability a model assigns to stereotypical versus non-stereotypical sentence completions (e.g., “girls/boys are smart”). However, some of the items within such benchmarks are unlikely to actually measure relevant bias (e.g., “the civil servant was a bureaucrat/bee keeper at

¹⁴ This language is adapted from Hannah Quay-de la Vallee, Center for Democracy & Technology, *Sharing the Health Guidance for Schools When Procuring Mental Health Technologies* 26 (2023), <https://cdt.org/wp-content/uploads/2023/06/2023-06-28-Civic-Tech-Guidance-for-Schools-When-Procuring-Mental-Health-Technologies-report.pdf>

¹⁵ Ryan Burnell et al, *Rethink Reporting of Evaluation Results in AI*, *Science*, vol. 380, no. 6641, pp. 136-138 (Apr. 2023), <https://eprints.whiterose.ac.uk/198211/1/Burnell%20et%20al%20-%202023%20-%20Rethink%20reporting%20of%20evaluation%20results%20in%20AI%20%28Pre-print%29.pdf>.

heart, so he knew where he really belonged”).¹⁶ Instance-level evaluations can help agencies appropriately calibrate the applicability of the overall evaluation metrics for their specific use cases.

- **Interactive interfaces that allow agencies to perform example-based interrogations:** In many cases, agencies may have data that they would like to evaluate a system against, but that they are not able to provide to vendors before a contract is in place. In order to allow for this sort of evaluation, vendors should, ideally, provide an interface that allows agencies to see what system outputs result from their own inputs of interest. The OpenAI playground¹⁷ and the inference API available for some Hugging Face models¹⁸ are examples. Interactive interfaces can allow agencies to test relevant edge cases that are of particular concern and help them build up intuition about how the system works. Interactive interfaces also allow agencies to reproduce vendors’ evaluations and make it less likely that vendors can base their evaluations on cherry-picked examples where the performance is good.

6: Which elements of testing, evaluation, and impact assessments are best conducted by the vendor, and which responsibilities should remain with the agencies?

Vendors will be better positioned to conduct an initial assessment of the impact of the data used to train and test their models, but agencies should play a role in assessing whether their desired use cases sufficiently resemble the conditions of deployment to determine whether a vendor’s testing, evaluations, and impact assessments will be reliable signals of true performance and risk.¹⁹ Consequently, beyond initial testing, much of the evaluation and impact assessment process will need to be done by agencies (perhaps in partnership with vendors), because the performance and potential impact of a system are contextual and will depend on the deployment context. That context includes what data will be used by the system and how results of the system will be employed, which the vendor will likely not have sufficient access to before a contract is in place. Consequently, *The Federal Government’s Power of the Purse* notes that agencies should make broader use of their authority to conduct pre-award evaluations for

¹⁶ Su Lin Blodgett et al, *Stereotyping Norwegian Salmon: An Inventory of Pitfalls in Fairness Benchmark Datasets*, in Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing 1004, 1008-09 (2021), <https://aclanthology.org/2021.acl-long.81.pdf>.

¹⁷ OpenAI Playground, <https://platform.openai.com/playground>.

¹⁸ See e.g., Hugging Face, BERT Base Model Card, <https://huggingface.co/google-bert/bert-base-uncased>.

¹⁹ Usman Anwar et al, *Socially Responsible Language Modelling Research (SoLaR) Workshop*, NeurIPS (Dec. 16, 2023), <https://neurips.cc/virtual/2023/workshop/66526>.

AI models:²⁰ “Although agencies have the authority to conduct technical evaluations before contract award, many agencies do not use ‘show, not tell’ demonstrations or require rigorous, independent testing and evaluation strategies to eliminate vendors.

GAO has explained, and comments on the Proposed OMB AI Memo have reinforced, that government evaluation of AI should be iterative and that ‘to manage technical performance, AI technical stakeholders—data scientists, data engineers, developers, cybersecurity specialists, program managers, and others—will have to ensure that the AI system solves the problem initially identified; uses data sets appropriate for the problem; selects the most suitable algorithms; and evaluates and validates the system to ensure it is functioning as intended.’

If a vendor is unable to meet the government’s requirements, the government would be better off eliminating them pre-award. To accomplish this requires the government to conduct pre-award evaluations and technical demonstrations.”

In addition to vendor testing and their own evaluations, agencies should also incorporate independent auditing:²¹ “Agencies should consider creating a separate contract vehicle to independently audit and evaluate vendors’ AI performance. Similar to the use of Independent Verification & Validation (IV&V) services operated either generally through formal, centralized Program Management Offices or specifically as contract actions for individual projects or programs, agencies should consider establishing formal methods to retain independent reviewers of vendors’ AI claims.

Contracts should include regular reporting and auditing on system performance (including any emergent biases and efforts on bias mitigation), which would need to be included either as deliverables or explicit contract terms.”

Ultimately, pre-deployment testing, evaluation, and impact assessments should not be taken as a guarantee against harm or error after a system is deployed; vendors must be expected to monitor for adverse events and impacts, to report them to agencies when anomalous or erroneous behavior is detected, and to take reasonable efforts to remediate errors that have led to those impacts.

7: What, if any, terms should agencies include in contracts to protect the Federal Government's rights and access to its data, while maintaining protection of a vendor's intellectual property?

²⁰ See *The Federal Government’s Power of the Purse* at 37 (citations omitted).

²¹ See *The Federal Government’s Power of the Purse* at 39-40 (citations omitted).

The appendix of CDT's *Sharing the Health: Guidance for Schools When Procuring Mental Health Technologies* recommends the following contract terms (adapted here for broader use, as the original work addresses an education-specific context).²²

- **Data governance.** The product provided by the vendor must adhere to the following specifications:
 - The product must allow the agency to customize data that is collected, including the ability to minimize sensitive data fields from collection if that data is not required for the functioning of the product.
 - The product must allow the agency to customize pre-set retention timelines for all elements of individual data, after which time agency data will be deleted from the product and from the vendor's data storage system unless an authorized administrator requests that the data be retained.
 - The product must allow administrators to delete individual information upon request. The vendor must be able to certify that the data has been deleted using industry best-practice deletion methods. The vendor must provide a certificate of deletion signed by the technical leader for data management.
 - The product must include a role-based access management system controlling agency data. Those roles must include *[insert required roles]*.
 - The vendor may not use the data collected by and stored in the product for any purposes other than providing the contracted services.
 - The vendor shall notify the agency of any data breach, security breach, or suspicion of such a breach within *[insert time frame]*.
- **Strong privacy and security controls**
 - The agency shall retain control of all data created or used by the product. The data shall not be reused or repurposed for any use other than the delivery of the product services unless explicitly allowed by the agency.
 - The vendor shall use industry-standard best practices for the secure transmission and storage of all agency data.
 - The vendor shall provide the agency with primary points of contact for those responsible for the privacy and security of the agency's data.
 - The vendor will provide privacy and security training for all of their staff who will access or interact with agency data.

²² See *Sharing the Health: Guidance for Schools When Procuring Mental Health Technologies*, *supra* note 10, at 25-27.

- All security and privacy assurances provided by the vendor shall be memorialized in a formal contract or related data-sharing agreement.
- **Sub-vendor compliance.** If the vendor employs any sub-vendors for any component of their product or system, the vendor will certify that the sub-vendor will also comply with any relevant requirements set out for the primary vendor.

10. How might OMB ensure that agencies procure AI systems or services in a way that advances equitable outcomes and mitigates risks to privacy, civil rights, and civil liberties?

Part V of *The Federal Government's Power of the Purse* recommends that, when reviewing vendor contracts involving AI, agency counsel should review for consistency with “civil rights laws, and compliance with [Executive Order 14091]’s mandate for agencies to acquire and use AI in a manner that advances equity. Agency counsel may also review a vendor’s proposed means for testing for and remediating bias to ensure they are sufficient and comport with relevant legal requirements. Contracts that implicate personally identifiable information or government data should receive additional review and attention.”²³

Performance of both PIAs and AIAs during procurement should also include examining and mitigating risks to privacy, civil rights, and civil liberties. Part V of *The Federal Government's Power of the Purse* recommends the following:

- An AIA should include the system’s “intended purpose, expected benefits, potential risks of the AI use, []the quality and appropriateness of the data used...and both general and sector-specific impacts and details about non-AI approaches that were considered.”²⁴ OMB should require agencies “to publish their AIAs for each of their AI use cases in a manner that is easy to find through their AI inventory submissions,” with “sufficient detail and raw figures for the public to understand how the AI was trained, its outcomes for each affected demographic group, and how the agency measures the impacts on these groups.”²⁵
- NIST should adopt a standard for scoping AIAs according to the degree of impact. “Although NIST has adopted the AI RMF, it should go further to develop a standard for federal AIAs (akin to NIST FIPS 199 [Standards for Security Categorization of Federal Information and Information Systems] / SP 800-53 [Security and Privacy Controls for

²³ See *The Federal Government's Power of the Purse* at 34 (citations omitted).

²⁴ See *The Federal Government's Power of the Purse* at 47 (citations omitted).

²⁵ See *The Federal Government's Power of the Purse* at 47 (citations omitted).

Information Systems and Organizations]), ideally with different levels associated with the degree of impact (akin to the FIPS 199 low/moderate/high impact approach). Any such standard must incorporate input from all affected stakeholders, including civil rights groups and impacted communities.”²⁶

In addition to these frameworks, effectively assessing AI systems will require internal expertise at agencies. To that end, Part V of *The Federal Government’s Power of the Purse* recommends the following approaches to strengthen agencies’ AI capacity:²⁷

- **“Develop training modules that ideally would be incorporated into the existing procurement curriculum.** The 2023 AI EO requires each agency to implement (or expand existing) AI training programs for their workforce. Although standalone training modules could be helpful, they are more likely to affect federal procurement practices and policies if they are embedded in training opportunities and curricula that are already in use.” For example, as these comments discuss in the response to Question 1, the DITAP and Federal Acquisition Institute curricula could incorporate responsible AI training, including on protection of privacy, civil rights, and civil liberties in agencies’ procurement and use of AI.
- **“Encourage growth and support of digital services teams within the government with experience designing and deploying responsible AI.** A number of organizations within the government are focused on recruiting, developing, and retaining technical talent that can bring focused capacity to some of the larger challenges within the federal procurement ecosystem. For example, the US Digital Service, the GSA CoE, Senior Advisors for Delivery, and the Presidential Innovation Fellows programs all bring technical expertise to assist agencies in their missions. These programs would benefit from the inclusion of training and information on AI, including its responsible procurement.”

Conclusion

CDT appreciates OMB’s continued leadership on AI policy and commitment to ensuring that people’s rights and safety are protected from AI harms by advancing responsible AI procurement practices in the federal government. We strongly encourage OMB to adopt the recommendations offered in these comments and in *The Federal Government’s Power of the*

²⁶ See *The Federal Government’s Power of the Purse* at 48 (citations omitted).

²⁷ See *The Federal Government’s Power of the Purse* at 49-50 (citations omitted).



Purse when applying M-24-10 to federal procurement processes, and we look forward to supporting this work.

Sincerely,

Elizabeth Laird
Director, Equity in Civic Technology

Hannah Quay-de la Vallee
Senior Technologist, Equity in Civic Technology

Ridhi Shetty
Policy Counsel, Privacy & Data