MITRE's Response to the OMB RFI on AI Acquisition

April 29, 2024

For additional information about this response, please contact:

Duane Blackburn
Center for Data-Driven Policy
The MITRE Corporation
7596 Colshire Drive
McLean, VA 22102-7539

policy@mitre.org
(434) 964-5023

About MITRE

MITRE is a not-for-profit company that works in the public interest to tackle difficult problems that challenge the safety, stability, security, and well-being of our nation. We operate multiple federally funded research and development centers (FFRDCs), participate in public-private partnerships across national security and civilian agency missions, and maintain an independent technology research program in areas such as artificial intelligence, intuitive data science, quantum information science, health informatics, policy and economic expertise, trustworthy autonomy, cyber threat sharing, and cyber resilience. MITRE's ~10,000 employees work in the public interest to solve problems for a safer world, with scientific integrity being fundamental to our existence. We are prohibited from lobbying, do not develop or sell products, have no owners or shareholders, and do not compete with industry—allowing MITRE's efforts to be truly objective and data driven. Our multidisciplinary teams (including engineers, scientists, data analysts, organizational change specialists, policy professionals, and more) are thus free to dig into problems from all angles, with no political or commercial pressures to influence our decision making, technical findings, or policy recommendations.

MITRE has a 50-year history of being a key partner to federal agencies, applying artificial intelligence (AI) and machine learning (ML) to advance missions while ensuring ethical safeguards. Our experience spans the entire AI/ML lifecycle, equipping us to anticipate and resolve future needs crucial to public safety, well-being, and success.

MITRE's Cost, Acquisition, and Management Sciences (CAMS) Innovation Center provides multidisciplinary analyses and products, enabling sponsors to effectively and affordably acquire, create, and deploy systems. Leveraging our project work across defense, civil, and intelligence agencies, CAMS applies innovative acquisition strategies to solve significant challenges.

Responding to this Request for Information (RFI), we draw on our extensive AI and federal acquisition experience to provide insights. We acknowledge the criticality of AI procurement by federal agencies and the need for careful consideration and innovative approaches.

Introduction and Overarching Comments

Our analysis suggests that standard practices and strategies of federal procurement may need to be adapted to accommodate the unique challenges and needs of AI procurement. Treating AI procurement as any other procurement could lead to unnecessary roadblocks and delays. Instead, we suggest considering standard practices and strategies as starting points that can be modified to meet the specific requirements of AI procurement.

This response also highlights the potential benefits of expanding the use of Other Transaction (OT) authority to specifically address AI procurements. This could provide a means of attracting and partnering with Non-Traditional Vendors (NTVs) and ensuring maximum competition.

Based on our analysis, a specialized training and certification program for AI procurements could provide uniform, documented certifications for Acquisition Workforce Professionals (AWPs), supporting their education for current and future AI procurement success.

Technically, consider the capabilities required for demonstrating large AI models in government platforms like FedRAMP. Contract language should support vendor demonstrations, considering the significant funding implications. Additionally, we recommend considering an approach similar to the Software Bill of Materials for AI systems. This would provide transparency about the components used in AI systems, including the training data, which is crucial for understanding their operation, managing their risks, and ensuring their security.

Answers to Questions Posed in the RFI

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 between agency missions and technical requirements?

Federal agencies are poised to acquire AI capabilities both directly and indirectly, as AI will be integrated into a multitude of systems and tools procured for specific purposes. This dual approach amplifies the complexity of "AI acquisition." Crucial to navigating this complexity is a deep understanding of the alignment between agency needs and technical specifications in an acquisition. Tools such as the RAI Toolkit,[1] developed by the Defense Innovation Unit (DIU) and the DoD Chief Digital and Artificial Intelligence Office (CDAO), can significantly aid in this process by supporting user engagement and facilitating requirements development.

With this context in mind, the Office of Management and Budget (OMB) must consider and accept that "standard practices and strategies" may not work for planned AI procurements. AWPs at all levels should consider existing standards and practices as guides, not rules, to maximize innovative approaches that will require modifying and creating new practices and strategies regarding planned AI procurements.

One of the most significant mistakes in procuring AI would be to treat it like any other procurement, by simply adhering to the mindset of 'this is how we usually do things,' and consequently applying the same methods and constraints to planned AI procurements. This creates unnecessary roadblocks, time delays, and avoidable rework, and often results in loss of interest by viable vendors. The preferred alternative is accepting that AWPs supporting a planned AI procurement will be working from a blank slate and must treat standard practices and strategies as starting points, accepting they'll need to be modified where allowable to meet the unique challenges and needs of planned AI procurements.

MITRE recommends an expanded focus on use of OT authority and best commercial acquisition practices to specifically address AI procurements. This will be key to attracting and partnering with NTVs to ensure maximum competition and access to the best-in-class industry has to offer based on the government's diverse requirements across defense, intelligence, and civilian agencies. Use of OT authority is also ideal for AI, with uncharted risks such as Intellectual Property (IP) rights, data management and controls, ways to estimate price/cost and measure general performance, and need for security frameworks and response requirements. Federal Acquisition Regulations (FAR)–based procurements extremely limit the ability of government

[1] RAI Toolkit. 2023. CDAO, https://rai.tradewindai.com/. Last accessed April 12, 2024.

and industry to openly partner and collaborate to reach consensus on how to navigate these barriers to achieve shared success for planned AI procurements.

We also recommend expanding the GSA Acquisition Gateway to include informational documents regarding successful and failed attempts to procure AI. The AWPs, namely Federal Acquisition Certification in Contracting (FAC-C) and Federal Acquisition Certification for Contracting Officer's Representatives (FAC-CORs) (including Level 3's) do not have extensive experience planning for, structuring, and moving AI procurements through the Acquisition Lifecycle; therefore, it's critical that OMB account for the "Knowledge Management" components and ensure proper information, guides, tools, etc. are created and made available to AWPs as lessons learned so other agencies avoid repeating the same mistakes.

The Federal Acquisition Institute should also consider creating a specialized FAC-C, FAC-COR, and Federal Acquisition Certification for Program and Project Managers (FAC-P/PM) training and certification program for AI procurements. This would ensure uniformly established training and documented certifications for AWPs are available and achievable to support educating AWPs for current and future success in AI procurement, similar to past efforts delivering specialized training on procurement areas for Cloud, Category Management, Construction, Agile, and Cybersecurity.

There are numerous innovative acquisition practices that can support the emerging practices in AI acquisition, including the following procedures, innovations, approaches, and strategies:

- Truly tapping into the industrial base through the utilization of various Innovative Market Research and Industry Engagement techniques[2]
- Executing the Challenge-Based Acquisition approach and using operational and evaluated demonstrations in the source selection phase[3]
- Utilizing a Statement of Objectives (SOO) vs. a Statement of Work (SOW) and tapping into the power and innovation of industry versus being overly prescriptive through the use of a SOW
- Establishing an AI Center of Acquisition Excellence and/or AI Consortium/Community of Practice that shares use cases, best practices, lessons learned, and templates

The acquisition culture must be addressed as well with regards to acquiring AI. Again, there should be a place for Contracting Officers (COs) and Contract Specialists (CSs) to learn about AI in general. Eventually, policy/procedure and the FAR should incorporate best practices and begin institutionalizing them. Acquisition teams should be encouraged and provided with opportunities such as professional development programs and recognition initiatives to learn and develop best practices for acquiring AI, and to share them across the federal ecosystem.

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

[2] Level-Up Your Market Research Game: Strategies and Hacks You'll Want to Use. 2022. MITRE, https://www.mitre.org/sites/default/files/2022-04/pr-21-4001-level-up-your-market-research-game-strategies-hacks-2022-guidebook.pdf.

[3] S. Roe, et al. Challenge-Based Acquisition: 5th edition. 2020. MITRE, https://www.mitre.org/sites/default/files/2021-11/prs-20-0745-challenge-based-acquisition-version-5.pdf.

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

Promoting robust competition within the federal AI marketplace is a complex task, with challenges ranging from administrative burdens and data rights requirements to the alignment of business structures with federal processes. These issues are particularly pronounced for leading AI vendors and small AI innovators, which often find the pace of acquisition and budget lifecycles out of sync with their rapid innovation. Furthermore, small AI vendors may be hesitant to partner with large primes or big tech companies due to concerns over IP and innovation constraints.

To address these multifaceted challenges, a comprehensive approach is needed. This could include expanding OT authorities, similar to those granted to DoD under 10 U.S.C. § 4022 – "Authority of the Department of Defense to carry out certain prototype projects." Such an expansion would allow for the negotiation of customized AI solutions in a proven acquisition environment, free from the constraints of the FAR.

In parallel, the government can promote awareness of AI needs to potential AI cloud service providers, which must invest time and money in obtaining provisional Authority to Operate (ATO) for the FedRAMP marketplace. Sponsorship and support for small and medium businesses will attract new entrants and provide a larger pool of AI options. Awareness of the demand for AI capabilities must also be communicated to FedRAMP to ensure prompt attention for AI assessment packages. Given the practical aspects of DoD's concept of adding specific security requirements to FedRAMP assessments, enabling small, medium, and large AI cloud service providers to meet these requirements provides benefits to civilian, defense, and intelligence agencies.

Learning from successful models like DoD's CDAO Tradewinds procurement model, which enables a 30-90 day procurement lead time, could be beneficial. However, as both OTA and Tradewinds are DoD-specific, there's a need for similar solutions for the broader government. OMB should review and consider expanding the Tradewinds approach and prototype contracting capabilities government-wide. Consortia may offer an alternative beyond DoD, with agile management and digital innovation being key.

Data rights and IP rights, while interconnected, should be distinctly addressed. The federal government has the right to own and control all the data it supplies to AI services, as well as the intermediate and final results. This includes any IP the government provides in using the AI service.

On the other hand, the AI algorithm's IP, developed by the AI service provider, belongs to the provider. It is crucial to articulate to stakeholders that the government seeks appropriate rights over the data it supplies and the results generated, which is vital for program success and long-term viability. These rights can foster an open architecture for future competitions and technological advancements.

Clear delineation of data and IP rights can enhance competition during prototyping, allowing vendor selection based on value in terms of results and data rights. It is also important to avoid sole reliance on Contract Performance Assessment Reporting System–based past performance for selection, as this could potentially exclude most non-traditional entities.

3. Should the Federal Government standardize assessments for the benefits and trade-offs between in-house AI development, contracted AI development, licensing of AI-enabled software, and use of AI-enabled services? If so, how?

Federal acquisition teams face unique challenges, including specific capabilities, requirement maturity, and resources, complicating standardization between in-house development and AI outsourcing. High-consequence AI capabilities demand heightened scrutiny on responsibility and transparency, which may not be sufficiently addressed by licensed AI software. Therefore, a flexible, case-by-case approach could more effectively address diverse team needs and circumstances than a standardized one.

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?

Agencies should proactively check if AI is part of their solutions in Requests for Proposals (RFPs) and PMRs (Program Management Reviews), ensuring integration of appropriate testing and continuous monitoring, which is crucial for AI performance and reliability.

User-focused metrics are vital for AI trustworthiness, but may not align with common AI research metrics such as F1 score, precision, or recall. The best user-focused metric often needs to be part of the effort and may not be pre-determined. For instance, translating a goal to increase work process throughput into a specific AI metric may not be straightforward.

Soliciting industry and user feedback on these metrics is vital. This feedback aids in AI performance metrics and developing meaningful, user-centric performance measures, emphasizing a collaborative, iterative process in defining and refining AI performance metrics.

5. What access to … 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?

For acquisitions involving low-consequence commercial AI, where standard license clauses could be applicable (especially for solutions developed at the contractor's expense), the usual procurement procedures may suffice. However, for more complex or high-stakes AI acquisitions, it is recommended to specially negotiate data rights. This approach aims to strike a balance between the need for access to critical AI components and the associated costs. Such a negotiation would consider the unique aspects of each acquisition, including the nature of the AI system, the sensitivity of the data involved, and the potential impact of the AI system on agency operations. This tailored approach helps ensure that the agency retains necessary access and control while also respecting the rights and interests of the vendor.

In addition, vendors should supply sufficient information about their advanced analytic to answer the following:

- Information about training data and how it has been sourced. This is important, for example, to assess potential bias and copyright infringement. This stipulation is also

outlined in OMB's M-24-10 memorandum,[4] which emphasizes that data used to develop an advanced analytic often constitute a public asset.

- Sufficient information to allow scientists and regulators to determine what testing has been done for safety. This stipulation is highlighted in several places in OMB's M-24-10 memorandum.[5]

The government should require data cards as specified in OMB Memo M-24-10.[6]

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?

In addressing the question of what terms agencies should include in contracts to protect the federal government's rights and access to its data, while maintaining protection of a vendor's intellectual property, it's important to consider the evolving landscape of case law with AI and IP, such as the NYTimes vs. OpenAI case. The U.S. Government should remain agile and flexible in adapting to changing case law.

The Office of the Undersecretary of Defense for Acquisition and Sustainment (OUSD(A&S)) has developed guidance that could be instructive in this context. In June 2022, OUSD(A&S) published the Department of Defense Response to Responsible AI Strategy & Implementation (RAI S&I) Pathway. Subsequently, in August 2023, it published a series of recommended guidance presentations to implement this strategy.

Specifically, the RAI S&I Pathway Line of Effort 3.1.3 provides recommendations to "identify best practices related to strategies to preserve government IP in the acquisition of AI and AI-enabled systems." The aim is to ensure an open architecture of secure data deliverables and rights that support: protection of government IP, best-value determinations, avoidance of vendor lock, and oversight of DoD use of AI capabilities and adherence to DoD AI Ethical Principles.

This guidance addresses several key areas, including IP fundamentals, considerations for FAR versus Other Transaction Authority, recommended solutions and best practices for promoting competition, source selection, negotiations, avoiding vendor lock, and the development of an IP Strategy and Implementation Plan. It also includes best practices and recommendations from multiple agencies, such as DIU and CDAO.

This model, developed by DoD, should be taken into account as a valuable input when higher-level AI policies and guidance are being established. It's also important to ensure consistency with other guidance, such as FedRAMP, and to recognize the interest of legislative bodies in this area.

Additionally, in certain cases such as the acquisition of Generative AI software that will be executed against internet data, we recommend that the federal government carefully examine the

[4] OMB Memo M-24-10 on "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence", Section 4.d.ii, March 28, 2024

[5] OMB Memo M-24-10 on "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence", Sections 5.c.iv.B and 5.d.vii, March 28, 2024

[6] OMB Memo M-24-10 on "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence", Section 5.d.ii, March 28, 2024

degree to which it should hold vendors liable for gross misinformation generated as a result of the vendor's advanced analytic model (e.g., if a vendor improperly discloses model bias or model limitations).

8. What if any terms, including terms governing information-sharing among agencies, vendors, and the public, should be included in contracts for AI systems or services to implement the AI M-memo's provisions regarding notice and appeal (sections 5(c)(v)(D) and (E))?

When acquiring advanced analytic services, algorithms, and data, agencies should require vendors to clearly define use limitations, such as limitations on sharing of data. While we recommend that the federal government obtain clear IP rights, as outlined in our answer to question 7 (for example, of data produced once an analytic product has been procured), MITRE recommends clearly negotiating any sharing limitations, particularly with regard to data.

9. How might agencies structure their procurements to reduce the risk that an AI system or service they acquire may produce harmful or illegal content, such as fraudulent or deceptive content, or content that includes child sex abuse material or non-consensual intimate imagery?

In addition to continuous testing, agencies should incorporate specific requirements in their contracts that mandate vendors to have robust content moderation and filtering mechanisms in place. Furthermore, they should require vendors to demonstrate their ability to update and adapt these mechanisms as new forms of harmful or illegal content emerge.

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?

The nature of AI-based systems and services requires specialized methods and instruments for identifying and mitigating risks to privacy and civil liberties. This must start with relevant information requirements as an integral (and not just pro forma) part of the procurement process. In particular, the procurement process should explicitly require the provision of data and model "cards" (their specifications defined by OMB) so that information regarding the provenance and characteristics of AI models and data is available for assessment. Although M-24-10 highlights these for purposes of transparency and performance improvement, both types of information will be critical for meaningful assessment of privacy and civil liberties risks to individuals, which should be a *precondition* for the acquisition or acceptance of such systems and services.

These assessment methods and instruments must be fit for purpose, one aspect of which involves addressing systems' fitness for purpose, including the scientific validity of underlying assumptions. Note that unless appropriately and substantially enhanced beyond their currently limited focus, the privacy impact assessments (PIAs) performed under the E-Government Act will prove wholly inadequate for the evaluation of privacy risks in this context, not least because they concern themselves with only the end product and not the development process that led to

it. (However, nowadays they are also inadequate even for more conventional systems.[7]) Meaningful assessment of civil liberties risks likewise will require a method and instrument suitable for the task. While M-24-10 refers to AI impact assessments, experience with PIAs has amply demonstrated that, in the absence of detailed substantive guidance, such assessments will be constructed and executed in the most anodyne fashion possible. It is imperative that OMB leverage existing work on algorithmic impact assessments and issue guidance that is much more specific and risk-based than what is contained in M-24-10, even going so far as to design and promulgate a model assessment template and associated guidance.

Certain rights-impacting systems may very well merit additional measures to ensure trustworthiness. While M-24-10 notes that the required minimum practices for such systems are just that, realistically (and again reflective of experience with PIAs and privacy risk management) minimum practices are likely to also become maximum practices absent further direction from OMB regarding additional techniques and the criteria for employing them. For example, assurance cases based on structured argumentation have long been used to help ensure safety and have been adapted for security, privacy, and even ethics. There is no reason this technique could not be further extended to AI-based systems and services to address privacy as well as civil liberties and, indeed, safety concerns. Similarly, substantial research exists on algorithmic audits to confirm that privacy and civil liberties risk mitigations are effective and that new risks have not manifested. However, without more concrete guidance, the ongoing monitoring called for in M-24-10 is likely to be de minimis.

While M-24-10 emphasizes intra-department and agency coordination, it barely touches on inter-department and agency coordination and collaboration. While the interagency AI council called for in Executive Order 14110 is a worthwhile endeavor, it may struggle to usefully engage with all the associated implementation issues. OMB should actively encourage other such bodies, including the Federal Chief Information Officer and Privacy Councils, to take up some of these issues where relevant. In particular, these existing mechanisms could be leveraged to facilitate sharing and even development of best (as opposed to minimum) practices.

Bonus – Other topics with implications for the procurement of AI by Federal agencies.

Transparency about the origins and development of AI models, as well as the datasets used for their training and evolution (aka an AI's supply chain), is crucial for assuring the AI systems that the government procures, develops, or uses. A bill of material of these AI models and data sets enhances this transparency, providing assurance about the construction, training, operation, and use, as well as the supporting assurance cases.

The Linux Foundation/OMG SPDX 3.0 standard[8] for bills of materials includes defined profiles for AI and datasets, offering the necessary composition information. Similarly, the OMG Structured Assurance Case Metamodel[9] describes a standard method for sharing machine-readable Assurance Cases.

---

[7] MITRE's Response to the OMB RFI on Privacy Impact Assessments. 2024. MITRE, https://www.mitre.org/sites/default/files/2024-04/PR-23-02057-21-MITRE-Response-OMB-RFI-Privacy-Impact-Assessments.pdf.

[8] System Package Data Exchange (SPDX®) Version 3.0. 2024. MITRE, https://www.omg.org/cgi-bin/doc?admtf/24-03-17.pdf.

[9] Structured Assurance Case Metamodel. 2021. Object Management Group (OMG), https://www.omg.org/spec/SACM/2.2/About-SACM. Last accessed April 12, 2024.