




EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF MANAGEMENT AND BUDGET
WASHINGTON, D.C. 20503

THE DIRECTOR

September 24, 2024

M-24-18

MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

FROM: Shalanda D. Young 

SUBJECT: Advancing the Responsible Acquisition of Artificial Intelligence in Government

1. OVERVIEW

The use of artificial intelligence (AI) in the Federal Government presents tremendous opportunity for modernizing agency operations and improving the delivery of government services to the public, provided that the risks presented by the use of AI technology are mitigated. Realizing this goal involves recognizing that AI poses novel types of risk, and proactively integrating considerations for AI risk management into agency acquisition planning. This memorandum builds on previous efforts to harness the power and utility of AI in service of agency missions while protecting the public from potential risks or harms.

Consistent with the Advancing American AI Act (“the Act”),¹ Executive Order 14110, *Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*, and Office of Management and Budget (OMB) Memorandum M-24-10, *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence*,² this memorandum directs agencies to improve their capacity for the responsible acquisition of AI. It contains new requirements and guidance for agencies on establishing meaningful cross-functional and interagency collaboration to reflect new AI responsibilities, managing AI risk and performance, and promoting a competitive AI market with innovative acquisition. As required by the Act, interagency consultation was conducted throughout the development of this memorandum.

Ensuring Cross-functional and Interagency Collaboration. Cross-functional collaboration throughout the acquisition lifecycle has long been a foundational principle of the Government’s acquisition practices, and is no less critical for the acquisition of AI. This memorandum requires agencies to create or update acquisition policies, procedures, and practices to reflect new responsibilities and governance for AI, as established by OMB

¹ Pub. L. No. 117-263, div. G, title LXXII, subtitle B, § 7224(d)(1) (codified at 40 U.S.C. 11301 note), <https://www.congress.gov/117/plaws/publ263/PLAW-117publ263.pdf>.

² OMB Memorandum M-24-10, *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence* (March 28, 2024), <https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf>.

Memorandum M-24-10. This memorandum also requires agencies to share information on acquisition of AI across the executive branch.

Managing AI Risks and Performance. Agencies are subject to existing risk management requirements relevant to AI. This memorandum does not replace or supersede these requirements but rather supplements those requirements. As such, this memorandum directs or guides agencies, as appropriate, to adjust their acquisition policies and practices to account for the distinctive ways that AI systems and services are developed, trained, and deployed. In particular, this memorandum builds upon OMB Memorandum M-24-10 by establishing acquisition-related practices that agencies must implement to ensure effective deployment of required risk management practices for rights-impacting and safety-impacting AI. These include specific actions designed to address complex issues related to privacy, security, data ownership and rights, and interoperability³ that may arise in connection with the acquisition of an AI service or system. Additional practices are required or recommended to ensure responsible acquisition of generative AI and AI-enabled biometric systems.

Promoting a Competitive AI Market with Innovative Acquisition. The AI marketplace encompasses a wide range of providers who perform diverse tasks to facilitate AI adoption, including data collectors and labelers, model developers, infrastructure providers, system integrators, and AI service providers. A decision to acquire a particular service or system from a given vendor can influence or limit options for future acquisitions. To ensure agencies can always procure state-of-the-art AI, this memorandum requires agencies to prioritize careful decision-making for interoperability and take steps to prevent vendor lock-in. This memorandum also strongly encourages agencies to utilize innovative practices that can help agencies get the best outcomes from their AI acquisitions, and support a diverse, competitive, and resilient Federal marketplace for AI. Appendix I provides additional details on specific leading and innovative strategies agencies should consider, and steps to empower and enable the acquisition workforce with the training and skills necessary for acquiring AI effectively and responsibly.⁴

2. SCOPE

The Act, along with Section 10.1(d)(ii) of Executive Order 14110, directs OMB to develop an initial means by which to ensure that contracts for the acquisition of an AI system or

³ The term “interoperability” generally refers to the ability of two or more systems, products, or components to exchange information and use the information that has been exchanged, including to operate effectively together. This includes ensuring that open and standard data formats and application programming interfaces (APIs) are used so that foundational components can be used, including to build for new use cases, without obscure proprietary technologies or licensing.

⁴ Consistent with provisions of the Advancing American AI Act and Executive Order 14110 directing the publication of this initial means, this memorandum sets forth multiple independent requirements and recommendations for agencies, and OMB intends that these requirements and recommendations be treated as severable. For example, this memorandum’s provisions regarding the management of AI risks and performance in Section 4 are capable of operating independently, and serve an independent purpose, from provisions to ensure cross-functional and interagency collaboration in Section 3. Likewise, each of Section 4’s individual risk management practices serves an independent purpose and can function independently from the other risk management practices. Accordingly, while this memorandum governs only agencies’ own procurement of AI and does not create rights or obligations for contractors or the public, in the event that a court were to stay or enjoin application of a particular provision of this memorandum, or its application to a particular factual circumstance, OMB would intend that the remainder of the memorandum remain operative.

service align with the guidance for agencies provided in OMB Memorandum M-24-10 and advance the other aims identified in section 7224(d)(1) of the Act.

This memorandum is that initial means, and is scoped to address considerations associated with agencies' acquisition of an AI system or service, regardless of whether the acquired AI system or service is standalone or integrated into broader information technology (IT) products, offerings, or services. Given the varied nature of AI, this memorandum includes requirements for subcomponents of AI systems or services, such as requirements specific to models and data. The considerations addressed by this memorandum include "risks from the use of AI," as defined in Section 6 of OMB Memorandum M-24-10. This initial means does not address all considerations that may arise in connection with the acquisition of AI, such as those related to Federal information and information systems in general. This memorandum does not supersede other, more general Federal policies that apply to AI but are not limited in scope to AI, including policies relating to procurement, enterprise risk management, information resources management, competition, antitrust, data, privacy, accessibility, Federal statistical activities, or cybersecurity.

Agencies must continue to comply with applicable OMB policies relevant to AI, and to coordinate compliance across their components with all appropriate officials. Agency officials retain their existing authorities and responsibilities established in other laws and policies.

a. Covered Agencies. Except as specifically noted, this memorandum applies to all agencies defined in 44 U.S.C. § 3502(1).⁵ As noted in the relevant sections, some requirements in this memorandum apply only to Chief Financial Officers Act (CFO Act) agencies, as identified in 31 U.S.C. § 901(b). The requirements in this memorandum do not apply to elements of the Intelligence Community, as defined in 50 U.S.C. § 3003.⁶

b. Covered AI. This memorandum provides requirements and recommendations that, as described in more detail below, apply to AI systems or services that are acquired⁷ by or on behalf of covered agencies.

The term "AI system,"⁸ as used in the Act and this memorandum, includes data systems, software, applications, tools, or utilities "established primarily for the purpose of researching, developing, or implementing artificial intelligence technology," as well as data systems,

⁵ The term "agency," as used in both the AI in Government Act of 2020 and the Advancing American AI Act, is defined as "any executive department, military department, Government corporation, Government controlled corporation, or other establishment in the executive branch of the Government (including the Executive Office of the President), or any independent regulatory agency," but does not include the Government Accountability Office (GAO); the Federal Election Commission; the governments of the District of Columbia and of the territories and possessions of the United States, and their various subdivisions; or Government-owned contractor-operated facilities, including laboratories engaged in national defense research and production activities. 44 U.S.C. § 3502(1); see AI in Government Act of 2020 § 102(2) (defining "agency" by reference to § 3502); Advancing American AI Act § 7223(1) (same). As a result, independent regulatory agencies as defined in 44 U.S.C. § 3502(5), which were not included in the definitions of "agency" in Executive Order 13960 and Executive Order 14110, are covered by this memorandum.

⁶ <https://www.congress.gov/117/plaws/publ263/PLAW-117publ263.pdf>

⁷ This memo uses the term "acquisition" as defined by the Federal Acquisition Regulation (FAR), Subpart 2.101. The term "acquisition" thus includes "procurement."

⁸ A full definition of this term is provided in Section 6 of this memorandum.

software, applications, tools, or utilities where an AI capability “is integrated into another system or agency business process, operational activity, or technology system.” The term excludes, however, “any common commercial product within which artificial intelligence is embedded, such as a word processor or map navigation system.”

In determining whether a product that integrates AI functionality is excepted under this provision, agencies should assess both (1) whether the product is widely available to the public for commercial use, as opposed to products that are not readily available to the general public or are specialized or customized for agency use, and (2) whether the AI is embedded in a product that has substantial non-AI purposes or functionalities, as opposed to products for which AI is a primary purpose or functionality. For example, word processing software that is otherwise commercially available, but that the agency is acquiring in customized form to suit the agency’s needs, likely would be covered by the requirements of this memorandum, as it would not be a “common commercial product” within the meaning of the exception. Likewise, word processing software that is used primarily for its AI functionality likely would be covered by this memorandum. On the other hand, common commercial word processing software that has substantial non-AI purposes or functionalities, but for which AI is embedded for functions like suggesting text or correcting spelling and grammar, would likely fall within the exception and thus would not be covered by the requirements of this memorandum.

This memorandum does not govern:

- i. agencies’ regulatory actions designed to prescribe law or policy regarding non-agency uses of AI;
- ii. agencies’ evaluations of particular AI applications because the AI provider is the target or potential target of a regulatory enforcement, law enforcement, or national security action; or the agency is evaluating the AI application because it was used by a criminal suspect;⁹
- iii. agencies’ development of metrics, methods, and standards to test and measure AI, where such metrics, methods, and standards are for use by the general public or the government as a whole, rather than to test AI for a particular agency application;¹⁰
- iv. agencies’ acquisition of AI to carry out basic, applied, or experimental research¹¹ except where the purpose of such research is to develop particular AI applications within the agency; or
- v. AI used incidentally by a contractor during performance of a contract (e.g., AI used at the option of a contractor when not directed or required to fulfill requirements).

c. Contracts for Rights-Impacting or Safety-Impacting AI. OMB Memorandum M-24-10 requires agencies to implement minimum practices to manage risks for rights-impacting and safety-impacting AI by December 1, 2024, subject to certain extensions and waivers. In order to

⁹ AI is not in scope when it is the target or potential target of such an action, but it is in scope when the AI is used to *carry out* an enforcement action. For example, when evaluating an AI tool to determine whether it violates the law, the AI would not be in scope; if an agency was using that same AI tool to assess a different target, then the AI would be in scope.

¹⁰ Examples include agency actions to develop, for general use, standards or testing methodologies for evaluating or red-teaming AI capabilities.

¹¹ For more information about basic, applied, or experimental research, please reference the [National Science Foundation’s Frascati Manual](#). The full Frascati Manual and current and upcoming online Annexes are available at <http://oe.cd/frascati>.

implement OMB Memorandum M-24-10's requirements regarding rights-impacting and safety-impacting AI:

- i. By November 1st, 2024, agencies shall identify any contracts associated with agency use of rights-impacting or safety-impacting AI.
- ii. No later than December 1, 2024, agencies shall—
 - A. Ensure that any contracts identified as associated with agency use of rights-impacting or safety-impacting AI systems or services are brought into compliance with the requirements of OMB Memorandum M-24-10 and Sections 4(d), 4(e), and 4(f)(ii) of this memorandum.
 - B. Ensure that any new contracts issued in support of agency use of rights-impacting or safety-impacting AI are consistent with the requirements of OMB Memorandum M-24-10 and this memorandum.

d. Other Contracts for AI. This memorandum shall apply to any contract awarded pursuant to a solicitation issued on or after the date that is 180 days after issuance of this memorandum, as well as to any option to renew or extend the period of performance exercised on an existing contract after the date that is 180 days after issuance of this memorandum.

e. Applicability to National Security Systems. This memorandum does not apply to AI acquired for use as a component of a National Security System.¹²

As agencies implement this memorandum, OMB will work with members of the Federal Acquisition Regulatory (FAR) Council, the Chief Artificial Intelligence Officers (CAIO) Council, the Chief Acquisition Officers (CAO) Council, the Chief Information Officers (CIO) Council, the Federal Privacy Council (FPC), the cross-functional council established under Section 3(b)(ii) of this memorandum, and other interested stakeholders, as appropriate, to consider agency experiences, identify best practices, and determine the need and form of additional policy and procedural direction to support the efficient, effective, and responsible acquisition of AI.

3. ENSURING CROSS-FUNCTIONAL AND INTERAGENCY COLLABORATION

Agencies must engage in cross-functional collaboration, coordination, and information-sharing throughout the acquisition lifecycle for AI.¹³ This section addresses agency responsibilities to engage in enterprise-wide strategic planning and cross-functional transactional planning for acquiring AI. Agencies must ensure that existing management practices and procedures for acquiring IT are updated, or that new practices and procedures are established, as necessary, to include agency responsibilities in AI governance as detailed in OMB Memorandum M-24-10. At the enterprise level, agencies must effectively identify and prioritize AI investments that serve the agency's mission, develop capacity to deploy acquired AI, promote adoption of

¹² The term "National Security System" has the meaning provided in 44 U.S.C. § 3552(b)(6).

¹³ This is a foundational principle of acquisition, including IT acquisition. *See, e.g.,* FAR 7.104(a) and CAPITAL PROGRAMMING GUIDE – A-11, PART 7; OMB Circular A-130, *Managing Information as a Strategic Resource*, Appendix II. Additionally, such collaboration and consultation is necessary to implement certain requirements under the Federal Information Technology Acquisition Reform Act (FITARA), which prohibits agencies subject to the CFO Act, besides the Department of Defense, from "...enter[ing] into a contract or other agreement for information technology or information technology services, unless the contract or other agreement has been reviewed and approved by the Chief Information Officer of the agency..." 40 U.S.C. § 11319(b)(1)(C).

best practices, and support management of risk associated with the use of AI that is acquired.¹⁴ At the transactional level, cross-functional collaboration will support consistent implementation of AI acquisition requirements, and better ensure that requirements reflect mission needs. This section also addresses the importance of regular interagency collaboration and information sharing about AI acquisition to strengthen the marketplace over time by increasing predictability and standardizing expectations for vendors.

a. Internal Cross-Functional Coordination on AI Acquisition.

In general, acquisition of AI systems and services should be treated as a type of IT acquisition for purposes of applying existing guidance.¹⁵ Similarly, acquisitions for other types of services where performance will be accomplished through the use of AI will remain subject to existing guidance. However, existing internal policies, procedures, and practices must be updated as necessary to reflect new considerations for acquiring AI, as described below.

i. Formalize Cross-Functional Collaboration to Manage AI Performance and Risks.

Each agency must establish or update policies and procedures for internal agency collaboration to ensure that acquisition of an AI system or service will have the appropriate controls in place to comply with the requirements of this memorandum, and that the agency's use of the acquired AI will conform to OMB Memorandum M-24-10. Within 180 days of issuance of this memorandum, agency CAIOs must submit written notification to OMB identifying progress made toward implementing this requirement, and identifying any challenges encountered or best practices identified during implementation. These policies and procedures should facilitate the cross-functional collaboration necessary to achieve timely acquisition and proactive risk management. Agencies must address:

- A. How planned acquisitions that involve an AI system or service will be initially reviewed by relevant agency officials to determine whether additional practices for managing AI performance and risk, as delineated in Section 4, are necessary;
- B. How officials with AI expertise and relevant equities (e.g., acquisition (including competition advocates), IT, cybersecurity, privacy, civil rights and civil liberties, budgeting, data, legal, program evaluation) are included in decision-making and coordination processes associated with the acquisition;¹⁶ and
- C. Conditions under which reviews and decision-making must be escalated, and to whom, including for planned AI acquisitions, implementing performance and risk management practices, monitoring and post-award management, and decommissioning.

¹⁴ This is necessary to implement the requirements of OMB Memorandum M-24-10 (e.g., Section 3(b)(ii)(U), Section 4(a)(i)-(ix)).

¹⁵ Such guidance includes, but is not limited to, OMB Circular A-130, *Managing Information as a Strategic Resource*. Section 3(c) of OMB Memorandum M-24-10 requires CAIOs "be involved, at appropriate times, in broader agency-wide risk management bodies and processes, including in the development of the agency risk management strategy." This memorandum makes clear that this includes processes for planning, programming, and budgeting, including for the acquisition of AI.

¹⁶ In addition to requirements under OMB Memorandum M-24-10 and FITARA, OMB Circular A-130, *Managing Information as a Strategic Resource*, details numerous requirements for agency-wide planning and coordination. For example, this includes a number of requirements for involvement of agencies' privacy programs and SAOPs. See generally Appendix II, Section 5(c)-(d) in OMB Circular A-130, *Managing Information as a Strategic Resource*.

- ii. **Ensure Agency-Wide Strategic Planning and Resourcing.** Within 180 days of the issuance of this memorandum, each agency CAIO must submit to OMB a plan for ensuring that the CAIO coordinates on AI acquisition with the CAO,¹⁷ CIO, Chief Information Security Officer (CISO), CFO, Senior Agency Official for Privacy (SAOP), and other relevant officials (e.g., Chief Technology Officer, Chief Data Officer, Chief Competition Officer, and an official responsible for protecting civil liberties). This plan should serve as a companion to each agency's enterprise AI strategy, including information on strategic planning and budgetary requests for the acquisition of AI. Strategic planning and forecasted budgetary requests must account for costs associated with risk management, including assessing and mitigating potential impacts to privacy, civil rights, civil liberties, and safety. This coordination plan should be reflected with written documentation, including through agencies' annual budget submission and relevant strategy documents, such as the agency's AI strategy.

b. Interagency Collaboration on AI Acquisition.

The following actions are designed to promote interagency collaboration consistent with harmonization efforts called for by Section 4(e) of OMB Memorandum M-24-10.

- i. **Centralization of Interagency Information and Knowledge Sharing.**
The CAIO Council, in consultation with OMB, the CIO Council, the General Services Administration (GSA) and the AI Community of Practice (CoP), will identify information and artifacts on AI acquisition to be collected and made available to all executive branch agencies. At a minimum, the CAIO Council should consider information such as:
 - A. Examples and lessons learned from successful and unsuccessful attempts to acquire AI, including sample requirements and contract clauses and provisions, and issues discovered in procured AI, especially for generative AI and other models which are commonly used across agencies;
 - B. Templates, including on novel and innovative AI practices as discussed in Section 4(c)(i) and Section 6, mechanisms to monitor and manage risk, disclosures from vendors (e.g., model and dataset cards, reports about company policies and processes);
 - C. Best practices, guides, methodologies (e.g., for responsible AI acquisition; testing, evaluation, and continuous monitoring; data access and restrictions related to the appropriate control of Federal data; and applying modular contracting practices in the AI context); and
 - D. Resources for assessing benefits and trade-offs between in-house AI development, contracted AI development, and licensing of AI-enabled software.

Each agency CAIO, CIO, CISO, and CAO should work together to collect and prepare the information and artifacts identified by the CAIO Council. The information should, where practicable and appropriate, be collected in a manner that supports standardized analysis and integration into the Hi-Definition Intelligent Acquisition Data Environment

¹⁷ Where an agency does not have a CAO, but has a Senior Procurement Executive (SPE), the CAIO should regularly coordinate and collaborate with the SPE.

(Hi-Def) , as aligned with OMB Circular A-137, *Strategic Management of Acquisition Data and Information*. Where possible, as determined by the CAIO, information and resources should be shared publicly to provide clarity to vendors, including new entrants.¹⁸

GSA, in collaboration with OMB and relevant interagency councils, should explore continued facilitation of the sharing of information, knowledge, and resources about AI acquisition across agencies, including through making any tools, resources, and data-sharing best practices developed for improved AI acquisition by relevant interagency council(s) or other relevant sources, available online to executive branch agencies, such as through a web-based repository. These efforts help to further the policy goals of OMB Circular A-137 that agencies collect and share acquisition data to support the mission of the Federal Government.¹⁹

- ii. **Cross-Council Working Group.** The Federal CIO and the Administrator for Federal Procurement Policy, in coordination with leaders of other executive branch interagency councils, will jointly establish a Federal cross-council working group, or identify an existing working group within an executive branch interagency council, to examine cross-functional issues that consistently arise in the procurement of AI. The working group will share the results with the IT Category Manager and GSA’s IT Vendor Management Office. This work will inform future acquisition strategies that can achieve a more diverse and resilient Federal supplier base for AI and can inform agency decisions to develop, buy, or build upon existing Federal AI systems and services.

4. **MANAGE AI RISKS AND PERFORMANCE**

The complex nature of how AI systems are built, trained, and deployed introduces additional considerations when agencies are acquiring AI from external entities. Agencies must ensure that relevant equities and risks are proactively considered when planning for an AI acquisition. Agencies should prioritize, at a minimum, privacy, security, data ownership, and interoperability as identified in Sections 4(b) and 4(c) of this memorandum.

Agencies must ensure their AI acquisitions comply with the risk management requirements identified in OMB Memorandum M-24-10 if the AI is used in a way that impacts rights or safety, while also continuing to prioritize privacy, security, data ownership, and interoperability. In such cases, agencies are required to implement the acquisition practices identified in Sections 4(d), 4(e), and 4(f)(ii) of this memorandum. Additional requirements are included below for the acquisition of generative AI and AI-enabled biometric systems. More specific advice and best practices for managing AI risks are available in the National Institute of

¹⁸ For a definition of “new entrant,” see OMB Memorandum M-23-11, *Creating a More Diverse and Resilient Federal Marketplace through Increased Participation of New and Recent Entrants*, <https://www.whitehouse.gov/wp-content/uploads/2023/02/M-23-11-Creating-a-More-Diverse-and-Resilient-Federal-Marketplace.pdf>.

¹⁹ OMB Circular A-137 establishes the principle that agencies should no longer view acquisition data as a singular agency asset, but rather an asset critical to supporting the missions of the Federal Government at large. As such, agencies should be prepared to collect and share acquisition data accordingly to support a hi-def environment where relevant information is available to the workforce at the point of need. <https://www.whitehouse.gov/wp-content/uploads/2024/05/OMB-Circular-A-137-Strategic-Management-of-Acquisition-Data-and-Information.pdf>

Standards and Technology's (NIST) AI Risk Management Framework (RMF) and the Blueprint for an AI Bill of Rights.

a. Determining Whether AI is Included in an Acquisition.

Agencies are required by Section 3 of OMB Memorandum M-24-10 to maintain and annually update an AI use case inventory. Understanding when AI is being acquired is also a prerequisite to managing the risks and performance of AI systems and services. To help agencies identify the acquisition of AI covered by this memorandum, officials responsible for acquisition planning, requirements development, and proposal evaluation should:

- i. Communicate to the vendor, to the greatest extent practicable, whether the acquired AI system or service is intended to be used in a manner that could impact rights or safety. In cases where an agency intends to procure AI capacity without full awareness of potential future use cases, the agency should decide during acquisition planning whether or not to require that any awards support use cases involving rights-impacting or safety-impacting AI, and plan accordingly;
- ii. In cases where an agency's solicitation does not explicitly ask for an AI system, consider requirements language asking vendors to report any proposed use of AI as part of their proposal submissions;
- iii. Require contractors to provide a notification to relevant agency stakeholders prior to the integration of new AI features or components into systems and services being delivered under contract. When notified, agencies should leverage their standard processes for determining whether risks from the use of AI are sufficiently managed, consistent with the requirements of OMB Memorandum M-24-10 and this memorandum, prior to accepting the contractor's proposed integration. This includes cases where integration of new AI features or components could impact rights or safety; in such cases agencies must ensure compliance with all applicable requirements for use of such AI; and
- iv. Communicate with vendors to determine when AI is a primary feature or component in an acquired system or service. This should also include questions to the vendor to understand if AI is being used in the evaluation or performance of a contract that does not explicitly involve AI.²⁰

b. Protecting Privacy, Civil Liberties, and Civil Rights.

- i. **Address Privacy Risks Throughout the Acquisition Lifecycle.** Given the increased privacy risks arising from the development, training, acquisition, and use of AI, agencies shall ensure that SAOPs and agency privacy programs have early and ongoing involvement in AI acquisition processes so that they are able to identify and manage privacy risks that may arise throughout the acquisition lifecycle of AI systems and

²⁰ For example, a vendor that has a contract to provide a report with recommendations on how to prioritize road construction based on current road conditions might use an AI model to forecast and assess traffic patterns for each neighborhood to inform recommendations. Understanding when vendors use AI in performance of a contract is particularly important where the vendors' use would qualify as rights-impacting or safety-impacting.

services and ensure compliance with law and policy related to privacy. Acquisition teams, such as Integrated Program Teams (IPTs),²¹ and other mechanisms for cross-functional collaboration shall include privacy expertise, consistent with OMB Circular A-130. Furthermore:

- A. Agencies shall ensure that the SAOP and agency privacy program are involved during pre-solicitation acquisition planning and while the agency is defining the requirements for prospective AI systems or services;
 - B. Consistent with OMB Circular A-130, agencies shall equip their privacy programs with the necessary resources to develop and implement, in coordination with other agency officials responsible for AI acquisition, appropriate contractual terms and conditions, policies, and processes to ensure that agency AI contracts comply with the privacy requirements in law and policy when they involve creating, collecting, using, processing, storing, maintaining, disseminating, disclosing, or disposing of information on behalf of a Federal agency or operating or using information systems on behalf of a Federal agency; and
 - C. Agencies should consider including as part of their evaluation criteria how AI vendors demonstrate they are protecting personally identifiable information (PII) and mitigating privacy risks, including through privacy-enhancing technologies.
- ii. **Ensure That AI-based Biometrics Protect the Public's Rights, Safety, and Privacy.** Expanding on Section 5(d)(vi) of OMB Memorandum M-24-10, agencies should ensure that contractual requirements address risks inherent in the procurement of AI systems that identify individuals using biometric identifiers (e.g., faces, irises, fingerprints, or gait), including risks that such an AI system is trained on or otherwise makes use of biometric data that was not lawfully collected or is not sufficiently accurate to support reliable biometric identification.²² To help address these risks for biometric identification and verification, agencies should avoid biometric systems that rely on unreliable or unlawfully collected information and ensure contractual terms address:
- A. Verification that AI-based biometric systems are not trained on data collected in violation of applicable law or Federal policy, and that such systems are sufficiently accurate to support reliable biometric identification and verification across different groups based on the results of testing and evaluation in operational contexts;
 - B. Requirements for vendors to submit systems that use facial recognition for evaluation by NIST as part of the Face Recognition Technology Evaluation and Facial Analytics Technical Evaluation, where practicable;
 - C. Requirements for supporting documentation or test results, as well as underlying test data where appropriate, sufficient to independently validate the operational

²¹ IPTs are designed to include appropriate cross-functional expertise (e.g., acquisition, IT, cybersecurity, privacy, civil rights and civil liberties, budgeting, data, legal, program evaluation). For more information on Specialized IT Acquisition Cadres and Integrated Product Teams, see generally, Memorandum from the Admin. for Fed. Procurement Policy & U.S. Chief Inf. Officer, *Acquisition Innovation Labs & Pilot for Digital Acquisition Innovation Lab* (Mar. 9, 2016), <https://techfarhub.usds.gov/assets/files/AIA-March9Memo.pdf>, and GAO-17-8, *IT WORKFORCE: Key Practices Help Ensure Strong Integrated Program Teams; Selected Departments Need to Assess Skill Gaps* (Nov. 2016), <https://www.gao.gov/assets/d178.pdf>.

²² This could include any such AI trained or operated using biometric data that embeds unwanted bias or was collected substantially without appropriate informed consent, or for another purpose, or without validation of the included identities.

performance of the AI system’s ability to match identities and the appropriateness of the relevant data used for training, through evaluations such as those offered by NIST or the testing facilities of the Department of Homeland Security’s Science and Technology Directorate. These independent assessments and benchmarks of systems should first include lab testing of algorithms, followed by operational testing, which should be continuously conducted using the operational version of the biometric system and be measured using standardized methodologies in as close to an operational context as possible; and

D. Requirements that biometric systems have certain properties:

1. Using a configurable minimum similarity threshold for candidate results;
2. Enforcing minimum quality criteria for input biometric data or samples used for biometric systems, including data used in reference galleries or training datasets for the system. These criteria should be based on standards set by independent bodies, such as the International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC)’s 29794-5:201;
3. For a given probe image or data input in use cases using a one-to-many identification search, returning a list of candidate matches above the minimum similarity threshold alongside similarity scores whenever possible; and
4. Maintaining detailed logs of use for auditing and compliance, including capturing input and output data in ways that incorporate appropriate protections for PII and other data throughout the information life cycle, and limiting retention and restricting reuse of PII for other purposes; system configuration parameters; resulting candidate matches, scores, and ordering; and other information as appropriate.

iii. **Comply with Civil Rights Laws to Avoid Unlawful Bias, Unlawful Discrimination, and Harmful Outcomes.** Many AI systems rely on vast amounts of data. These tools have the potential to produce outcomes that result in unlawful discrimination. Discrimination may come from different sources, including problems with data model opacity and access, and with system design and use.²³ Consistent with the risk management requirements of OMB Memorandum M-24-10, agencies should address risks that procured AI may generate unlawful bias, unlawful discrimination, or harmful outcomes, and require vendors to identify potential AI biases and mitigation strategies to address biases.

c. Practices for Managing Performance and Risk for Acquired AI.

i. **Use Performance-Based Acquisition Techniques that Enable Proactive Risk Management.** When acquiring AI, agencies should leverage performance-based approaches and techniques, including through the use of best practices delineated in Appendix I(a) of this memorandum, to strengthen their ability to effectively plan for, identify, and manage risk. Performance-based requirements allow agencies to understand and evaluate vendor claims about their proposed use of AI systems or services prior to contract award, acquire AI capabilities that address their needs, and perform post-award monitoring. Focusing acquisition on achieving desired performance outcomes directly

²³ Joint Statement on Enforcement of Civil Rights, Fair Competition, Consumer Protection, and Equal Opportunity Laws in Automated Systems (justice.gov)

facilitates an agency's ability to ensure agency needs are met by defining metrics to maintain and improve performance of the AI system or service.

- ii. **Ensure Performance Justifies Use.** The performance-based requirements identified by an agency, including safeguards against inaccurate outputs, confabulations, drift, and other risks specific to AI,²⁴ must ensure, to the greatest extent practicable, that the AI system or service to be acquired will be appropriate for the expected use contexts.
- iii. **Determine Appropriate Intellectual Property Rights and Ownership.** Consistent with applicable laws and governmentwide policy, an agency must include appropriate contractual terms that clearly delineate the respective ownership and intellectual property (IP) rights of the Government and the contractor. Careful consideration of respective IP licensing rights is even more important when an agency procures an AI system or service, including where agency information is used to train, fine-tune, and develop the AI system.

To that end, agencies must develop an approach to IP that considers what rights and deliverables are necessary for the agency to successfully accomplish its mission, protects Federal information used by vendors in the development and operation of AI systems and services for the Federal Government, considers the exploration of open-source development practices of AI code, avoids vendor lock-in, and avoids unnecessary costs. Agencies must scrutinize terms of service and licensing terms, including those that specify what information, models, and transformed agency data should be provided as deliverables, to ensure that they clearly articulate the scope of rights needed by the Government over its own data and any derived products.²⁵ Furthermore, agencies should conduct careful due diligence to the supply chain of a vendor's data. Best practices include the following:

- A. Negotiating the appropriate scope of licensing rights and other rights that are necessary to accomplish the Government's mission in the long term while avoiding vendor lock-in. This includes strategically selecting the appropriate FAR or agency supplemental clauses and making affirmative decisions about which alternates to these clauses are necessary. For example, as part of its acquisition planning, an agency may determine it needs unlimited rights to certain contractor deliverables based on its long-term approach to IP. Another agency may determine it requires assignment of copyright to deliverables specified in the contract. In all circumstances, agencies must consider its mission, long-term needs, and larger enterprise architecture while avoiding vendor lock-in and maximizing competition;
- B. Ensuring the contract clearly defines the process and timeline for delivery of components needed to operate and monitor the AI system, including as appropriate: data; inputs to the development, testing, and operation process;

²⁴ In identifying such risks and considerations, agencies may find it helpful to refer to resources and policies from other agencies, such as NIST's Generative AI Profile, <https://airc.nist.gov/docs/NIST.AI.600-1.GenAI-Profile.ipd.pdf>.

²⁵ As used here, the term "derived product" means any asset created during the process of transforming a data asset (e.g., modifying or enriching through any contractor-performed cleaning, labeling, customization, etc.) for the purposes of training, testing, or validating a model, as well as any AI models, systems, or services developed using a data asset. The term "data asset" has the meaning provided in 44 U.S.C. § 3502.

models; software; other technical components; and documentation as described in the agency's technical requirements. Contracts should ensure such components and their foundational code remain available for the acquiring agency to access and use for as long as it may be necessary (e.g., to re-train the model);

- C. Ensuring complete and timely delivery of information necessary to fulfill requirements of OMB Memorandum M-24-10, including incident reporting. This information should be provided in a machine-readable and native format for ingestion and analysis;
 - D. Requiring appropriate handling, access, and use of agency information, such as original input data, prompts, processed data, output data, weights, and models, at least in part by providing clear parameters to ensure that such information must only be collected and retained by a vendor when reasonably necessary to serve the intended purposes of the contract; and
 - E. Opting out of or prohibiting the contractor from using agency data to train AI without an agency's consent. The contract should permanently prohibit the use of inputted agency data and outputted results to further train publicly or commercially available AI algorithms, including generative AI, consistent with applicable law.
- iv. **Data Management.** The role of data as a strategic asset upon which AI systems are built makes the handling of that data a key contractual consideration for responsible AI acquisition. The terms of contracts must explicitly address how a vendor is to ensure (e.g., through a quality management system) compliance with relevant directives and policies, particularly with regard to:
- A. Systems and procedures for data management, such as data collection, data analysis, data labelling, data storage, data filtration, data mining, data aggregation, data retention, use, and any other operation regarding the data that is generated before, during, or after the delivery of the AI;
 - B. An accountability framework identifying the tiered levels of access and requisite responsibilities of handling data; and
 - C. Disclosures when copyrighted materials are used in the training data. Agencies may also consider whether similar disclosure of the use of synthetic or third-party data is useful within the context of specific acquisitions.
- v. **Approvals for Cybersecurity and Appropriate Protection of Agency Data.** In procuring AI systems and services, agencies are responsible for complying with existing law and governmentwide directives and policy related to cybersecurity and data security. In consideration of specific software controls related to data management and model security addressed in OMB Memorandum M-24-10, agencies should ensure the inclusion of contractual requirements that facilitate the ability to obtain any documentation and access necessary to understand how a model was trained, and to understand the integrity of models they intend to deploy for a specific use case. This can be achieved by:
- A. **Software Controls to Secure Training.** Agencies may consider requesting training logs from a contractor, to include evidence of any data sourcing, cleansing, inputs, parameters, or hyper-parameters used during training sessions for models delivered to the government, such that the training can be reproduced or verified if needed. When training occurs using agency data, agencies should look to existing data security software controls to ensure provenance and security.

Agencies may also wish to invest in infrastructure to evaluate software controls throughout training processes and data sources for projects at appropriate scale. Agencies may consider and adapt NIST guidance, such as Special Publication (SP) 800-218, *Secure Software Development Framework (SSDF) Version 1.1: Recommendations for Mitigating the Risk of Software Vulnerabilities*, to help guide these practices.²⁶

B. Software Controls to Secure Models. Contract requirements should address vendor compliance with all relevant agency requirements on protecting data, including software controls for privacy and security. Where practicable, vendors should provide detailed documentation of the training procedure used for the model to demonstrate the model’s authenticity, provenance, and security. Trained model artifacts should also be available for agency evaluation and review.

C. Evaluating the Model Design and Any Data Used in Model Training. Agencies should work to ensure contract requirements language establishes baseline expectations of model training using agency data to better understand how a model is expected to iterate using agency data. This is especially critical for capturing additional risks unique to AI, such as data poisoning or data leakage, which might occur when training a model with agency data inputs.

vi. **Ongoing Cost Management.** Agencies must leverage early-stage budgeting and financial planning (e.g., in cost estimates) to estimate costs for the ongoing operations and maintenance of AI systems, including post-award oversight, risk management, maintenance, and corrective action to ensure the AI system continues to operate appropriately and as intended.²⁷

d. Practices for Managing Risk and Performance for Rights-Impacting AI and Safety-Impacting AI.

As noted earlier in this memorandum, for existing contracts for AI systems or services where agency use is rights-impacting or safety-impacting, agencies must update contract terms as needed to comply with applicable requirements from Section 5(c) of OMB Memorandum M-24-10. This includes updates to contractual terms where an AI system or service was initially acquired for a use that did not impact rights or safety, but subsequently is used or planned to be used in a manner that does impact rights or safety. Agencies must cease use of AI systems or services that impact rights or safety in cases where required risk management practices cannot be sufficiently implemented,²⁸ as determined by the agency.

i. **Identify When Solicitations Require Compliance for Rights-Impacting and Safety-Impacting AI.** Where practicable, agencies must disclose in solicitations whether the planned use is rights-impacting or safety-impacting. In cases where an agency intends to procure AI capacity without full awareness of potential future use cases, the agency

²⁶ For example, SP 800-218A provides additional practices to guide security of generative AI and dual-use foundation models.

²⁷ “AI systems may require more frequent maintenance and triggers for conducting corrective maintenance due to data, model, or concept drift.” NIST AI RMF, Appendix B: How AI Risks Differ from Traditional Software Risks. This may result in more significant “Operations and Maintenance” budgets for AI systems.

²⁸ Unless a waiver, properly issued pursuant to the requirements of OMB Memorandum M-24-10, applies.

should decide during acquisition planning whether to require that any awards support use cases involving rights impacting or safety-impacting AI, and plan accordingly.

- ii. **Incorporate Transparency Requirements into Contractual Terms and Solicitations to Obtain Necessary Information and Access.** Agencies must ensure that vendors provide them with the information and documentation necessary to monitor the performance of an AI system or service and implement applicable requirements of OMB Memorandum M-24-10.²⁹ This may include information about the AI’s functionality and use that may be publicly posted in the agency’s AI use case inventory.

The level of transparency agencies must require of a vendor, both in the solicitation and evaluation process and through resulting contractual obligations, should be commensurate with the risk and impact of the use case for which the AI system or service will be used. Furthermore, careful consideration should be given to the range of potential agency use cases for the acquired AI system or service, and how the information required to facilitate compliance may depend on whether vendors are developers or deployers of an AI system or service. Agencies must consider whether any or all of the following categories of information must be provided by the vendor to satisfy the requirements of OMB Memorandum M-24-10 or to meet the agency’s objectives:

- A. Performance metrics, including real-world performance for specific sub-groups and demographic groups to surface discriminatory outcomes;
- B. Information about the training data, including the source, provenance, selection, quality, and appropriateness and fitness-for-purpose of the training data, the input features used, time period across which training data was collected, and any filters used;
- C. Information about programmatic evaluations of the AI system or service, including the methodology, design, data, and results of how the evaluation of the program delivering the AI system or service was conducted.;
- D. Information about testing and validation data, including the source, provenance, quality, and appropriateness and fitness-for-purpose of the testing and validation data, the time period across which it was collected, and the extent of overlap or other possible lack of independence from training data;
- E. Information about how input data is used, transformed, and retained by the AI and whether such data is accessible to the vendor;
- F. Information about the AI model(s) integrated into an AI system or service, including the model’s version, capabilities, and mitigations, to the extent it is available to the vendor;
- G. The intended purpose of the AI system or service, known or likely unintended consequences that may occur when deployed for the intended purpose, and known limitations; and
- H. Data protection metrics or assurance indicators for data in transit and at rest in AI systems.

²⁹ For example, information about the acquired AI system or service (e.g., its performance, its intended use) and relevant data (e.g., its quality and representativeness, how it was collected and prepared) are necessary for agencies to complete and updating impact assessments, independently evaluate the AI, and regularly evaluate risks from the use of the AI.

If the agency must obtain any of this information to satisfy legal or policy requirements, then the agency must incorporate requirements for the submission of that information into solicitation and/or contract documents.

- iii. **Delineate Responsibilities for Ongoing Testing and Monitoring and Build Evaluations into Vendor Contract Performance.** OMB Memorandum M-24-10 generally requires agencies to institute ongoing procedures to monitor degradation of the functionality of AI systems or services and to detect changes in their impact on rights and safety. However, there are instances when a vendor is best equipped to carry out those activities on the agency's behalf, and so is required under a contract to closely monitor and evaluate the performance and risks of an AI system. In such instances, agencies must still provide oversight and require sufficient information from a vendor to determine compliance with OMB Memorandum M-24-10. Agencies must ensure that contractual terms provide the ability to regularly monitor and evaluate (e.g., on a quarterly or bi-annual basis, based on the needs of the program) performance and risks throughout the duration of the contract. To do so:
- A. Agencies must use data defined by the agency (e.g., agency validation and testing datasets) when conducting independent evaluations to ensure the AI system or service is fit for purpose. To the extent practicable, the data used when conducting independent evaluations should not be accessible to the vendor, and should be as similar as possible to the data used when the system is deployed;
 - B. Contracts must require vendors to provide agencies with sufficient access and time to conduct any required testing in a real-world context, including testing carried out by others on behalf of or under agreement with the agency. Alternatively, agencies may require a vendor to regularly provide the results of an AI system or service's testing in a real-world operational context and the benchmarks used, with sufficient detail such that the testing could be independently verified or reproduced, if practicable;
 - C. Contracts must not prohibit agencies from disclosing how they conduct testing and the results of testing;
 - D. Contracts must detail the examination, testing, and validation procedures the vendor is responsible for and the frequency with which they need to be carried out;
 - E. Where appropriate, agency contracts for AI systems or services must also include terms that require vendors to provide the government with the results of performance testing for algorithmic discrimination, including demographic and bias testing, demographic characteristics of groups the performance testing has been conducted on, or third-party evaluations and assessments providing an equivalent level of detail. Alternatively, agencies may require a vendor to provide the results of performance testing to address these issues; and
 - F. Agencies must also consider how testing and monitoring, including as part of post-award management, impacts financial planning and budgeting requirements in Sections 3(a)(ii) and 4(c)(vi) of this memorandum.
- iv. **Set Criteria for Risk Mitigation and Prioritize Performance Improvement.** Agencies must have the ability, throughout the entire lifecycle of the contract, to update risk mitigation options and prioritize performance improvement of the AI system or service. To do so, agencies should consider:

- A. Contractual terms that require vendors to regularly monitor an AI system’s performance and rectify any unwanted system behavior, such as retraining the model or adding additional mitigations to the system, based on performance or event-based triggers. To the extent practicable, the cadence of reporting, review, and updating should be determined prior to award of the contract;
 - B. Contractual terms that require vendors to meet performance standards before deploying a new version of its AI system or service in performance of an agency contract or for a vendor to roll-back to a previous version if a new version fails to meet performance standards and requirements;
 - C. Incentivizing improved model performance through performance-based contracting (see Appendix I of this memorandum) and incentive contracts;³⁰
 - D. Contractual terms requiring vendors to participate in program evaluations sponsored by the agency to assess implementation and effectiveness, as an additional incentive to assess and improve the intended use case of an AI system, or service; and
 - E. Contractual language that requires vendors to document tools, techniques, coding methods, and testing results, as a means of promoting interoperability and mitigating vendor lock in.
- v. **Require AI Incident Reporting.** In addition to any existing reporting requirements for cybersecurity or other security-related incidents³¹ and breaches³² of PII, agencies must, to the greatest extent practicable, include contractual terms requiring that vendors have a process for identifying and disclosing to agencies serious AI incidents and malfunctions of the acquired AI system, or service within 72 hours, or a timely manner based on the severity of the incident, after the vendor reasonably believes the incident occurred. Criteria for what constitutes a serious AI incident or malfunction should be determined by each agency, and aligned with the vendor’s quality management system. For example, a serious AI incident or malfunction may consist of unexpected malfunctions, or unintended outcomes that directly result in harms to rights or safety, material or irreversible disruption of the management and operation of critical infrastructure, material damage to property, loss of a mission-critical system or equipment, failure of the mission of an agency, or loss of life.

Interagency collaboration described in Section 3(b) of this memorandum can support harmonized implementation of this requirement. Such implementation should include facilitating third-party and public discovery and reporting of AI incidents, as practicable and where appropriate.

e. Additional Practices for Managing Risk and Performance for Rights-Impacting AI Systems and Services.

³⁰ See FAR Subpart 16.4—Incentive Contracts

³¹ See, e.g., 44 U.S.C. § 3552(b)(2).

³² As defined in OMB Memorandum M-17-12, *Preparing for and Responding to a Breach of Personally Identifiable Information*, https://obamawhitehouse.archives.gov/sites/default/files/omb/memoranda/2017/m-17-12_0.pdf, § III(C).

For new or existing contracts involving agency use of rights-impacting AI systems or services, agencies must ensure these practices are followed for the acquired AI service or system:

- i. **Disclose OMB Memorandum M-24-10 Notice and Appeal Requirements to Vendors.** To the extent practicable and appropriate, agencies must identify in requirements documents for rights-impacting AI systems or services whether use of the AI will involve notifying individuals of AI-enabled decisions affecting them and affording opportunities for human consideration and remedy.³³ Agencies should consider whether to specify the information they will need to include in notifications to negatively affected individuals and to implement remedy processes, such as appeals.³⁴
- ii. **Require Vendors to Cooperate with OMB Memorandum M-24-10 Notice and Appeal Requirements.** Agencies should, to the greatest extent practicable, identify in their contractual requirements where vendor action is needed to support agency plans for notifying individuals when the use of AI results in an adverse decision, as described in OMB Memorandum M-24-10, or providing those individuals with opportunity to appeal. Contracts must include any requirements for additional access, information, or necessary documentation about the AI system or service necessary for the agency to carry out any plans for notice and appeal procedures. Agencies should also consider whether, to ensure successful implementation of a notice and appeal process, the contracts must specify the timeframes in which critical information will be provided by the vendor.

f. Additional Practices for Generative AI.

Consistent with Section 10.1(f)(i) of Executive Order 14110 and Section 4(b)(iv) of OMB Memorandum M-24-10, agencies are encouraged to provide their personnel and programs with access to secure and reliable generative AI, particularly for experimentation and carrying out low-risk productivity tasks. However, agencies must implement additional practices beyond the aforementioned risk management practices when acquiring general use enterprise-wide generative AI systems or services.

The term “general use enterprise-wide generative AI” refers in this memorandum to generative AI, in the form of a foundation model or other widely applicable generative AI system, that is acquired for general purposes for which the details are infeasible to define prior to procurement, such as for workforce-productivity use, general application development, or other general tasks, and is acquired for use by end users in more than one agency component (e.g., unit, division, office, bureau, agency, Service, Department); or through a contract vehicle that accommodates the requirements of more than one organizational component (e.g., component-wide contract, agency-wide contract, multi-agency contract, governmentwide acquisition contract, Multiple Award Schedules). The term does not include generative AI that is acquired to accomplish a specific or narrowly-scoped agency use case or that is only used to implement specific features of a system or service where such generative AI is not able to generate detailed responses to a wide array of open-ended inputs. This term does include access to generative AI provided or added to the offerings of a pre-existing contract or as a component of a larger

³³ See Sections 5(c)(v)(D) and (E) of OMB Memorandum M-24-10.

³⁴ Information might include the data that the decision relied upon, the design of the AI system, and the broader decision-making context in which the system operates as feasible.

contract, such as for an operating system or for cloud computing services, when the elements of the definition above are met.

The practices below³⁵ will help agencies address the unique challenges of such contracts, including to protect the rights and safety of the public, provide transparency about AI-generated outputs, and responsibly acquire generative AI when it is not feasible to know all the ways it may be used after it is acquired. Agencies should also monitor frameworks and standards published by NIST and appropriate international standards development organizations to update their approach to the acquisition of general use enterprise-wide generative AI accordingly. The practices for acquiring general use enterprise-wide generative AI include the following:

- i. **Provide Transparency About Risks and Generated Content.** When procuring general use enterprise-wide generative AI, agencies must include contractual requirements for vendors to:
 - A. Ensure that any audio, image, and video outputs of AI systems that are not readily distinguishable from reality are created or modified using mechanisms, such as through watermarks, cryptographically-signed metadata, or other technical artifacts,³⁶ that allow the outputs to be identified as generated by AI, attributed to the specific model that was used to produce the output, and linked with other relevant information about the origin or history of outputs;
 - B. Document how the general use enterprise-wide generative AI was or will be trained and evaluated, including relevant information about data, data labor, compute, model architecture, and relevant evaluations.³⁷
- ii. **Mitigate Inappropriate Use.** When procuring general use enterprise-wide generative AI, agencies must consider including contractual requirements that ensure vendors provide appropriate protections, where practicable, against the AI systems or services being used in ways that are contrary to law and policy. This may include providing methods for monitoring how the general use enterprise-wide generative AI is used in the agency and guidance for how to monitor such use effectively, as well as potentially implementing technical safeguards against the AI being used in prohibited or otherwise sensitive contexts, such as refusing prompts asking for prohibited outputs.

³⁵ These are best practices for contracting for generative AI. Inserting them into all contracts for generative AI may not be appropriate or feasible. Nevertheless, although this memorandum only mandates the inclusion of these requirements in general use enterprise-wide contracts for generative AI, agencies are strongly encouraged to seek the inclusion of these practices in all contracts for generative AI systems or services where appropriate, to the greatest extent practicable.

³⁶ For a survey of such mechanisms, refer to [NIST.AI.100-4.SyntheticContent.ipd.pdf](#). The effectiveness of these mechanisms continues to evolve both as they are improved and as mechanisms to bypass them improve. Agencies and their vendors should continue to assess the effectiveness of identification and labeling mechanisms and use those that are shown to be effective at providing transparency about AI-generated content and that are robust against relevant attacks and normal perturbations (e.g., cropping or resizing images).

³⁷ Agencies should determine relevant information, which may include data size, sources, selection, curation, augmentation, filtration; license status; the demographic distribution of the data; known or potential biases in the data; whether the data was collected with explicit consent for use to train or test AI; who collected and annotated the data; compute resources expended during training; duration of training; carbon emissions during training; input and output modalities of the model architecture; components of the model; model size; quantitative measurements of the model's capabilities, limitations, risks and mitigations, as practicable.

- iii. **Mitigate Risk of Harmful and Illegal Output.** When procuring general use enterprise-wide generative AI, agencies must include contractual requirements for vendors to:
 - A. Employ best practices to prevent the AI system from generating toxic, false, illegal, violent, or otherwise harmful content, including Child Sexual Abuse Material (CSAM) and Non-Consensual Intimate Imagery (NCII), such as by having appropriate policies, procedures, and technical measures to filter training data, prompts, and outputs;
 - B. Make best efforts to filter out known CSAM and NCII in their training data, such as by comparing their training data to the database of known CSAM maintained by the National Center on Missing and Exploited Children, where appropriate, as well as other similar known systems for NCII; and
 - C. Provide the agency with documentation on the components and capabilities in the model or system put in place by the vendor or configurable by agencies to limit harmful and illegal outputs.³⁸

- iv. **Testing Performance to Identify the Best Fit for Agency Mission.** When procuring general use enterprise-wide generative AI, agencies must take appropriate steps to provide an objective and empirical basis for competitively selecting the solution or solutions that provide the most value to the agency, including taking steps to compare the performance, cost, safety, security, and trustworthiness of available solutions under reasonably expected usage conditions.
 - A. Agencies must consider a range of alternatives to a proposed solution and must not privilege proprietary solutions over open-source alternatives solely on the basis of their proprietary nature, but instead make a decision based on comprehensive evaluation. It is important not to assume that proprietary solutions are inherently more secure or reliable than open-source alternatives.³⁹
 - B. To compare performance of general use enterprise-wide generative AI solutions prior to the procurement, agencies may estimate the most common general categories of beneficial tasks that the agency will use the AI system for, relying as appropriate on the Federal AI Use Case Inventory and any additional resources provided by OMB, GSA, or NIST. Agencies may then assess multiple available solutions based on their evaluated performance on these categories of tasks.

³⁸ This should include definitions of content filters and their severity thresholds.

³⁹ Securing generative AI models through methods such as opting to host open-source models can significantly enhance security by isolating systems from unsecured networks, but it also shifts the full burden of computational power and energy consumption onto the agency. This means that the agency must invest in and maintain the necessary hardware and infrastructure, leading to increased energy use and potentially higher environmental impact. Unlike developer-managed models, where the compute load is distributed across optimized infrastructure, open-source models may require agencies to handle all processing, which can be resource-intensive in the short-run, but might save significant costs in the long-run. Agencies must carefully weigh these trade-offs, considering whether the heightened security justifies the additional operational and environmental costs. While on-premise hosting of generative AI models may be necessary for protecting highly sensitive data, alternative approaches, such as secure cloud environments with strict controls or partnerships with developers offering dedicated infrastructure, could provide a more balanced solution when no wider enterprise AI infrastructure strategy exists. These alternatives might allow agencies to maintain robust security without overwhelming their resources or significantly increasing their carbon footprint.

- v. **Conduct Evaluations, Testing, and Red-Teaming and Share Results.** When procuring general use enterprise-wide generative AI, agencies must include contractual requirements ensuring that vendors:
- A. Provide documentation, including public documentation where feasible, on the following:
 1. Pre-deployment testing and evaluations, as well as testing and evaluations conducted on an ongoing basis;
 2. Red-teaming results, including results of any relevant testing or red-teaming by third-parties; and
 3. Steps taken to mitigate any issues discovered in the course of evaluations, testing, and red-teaming.
 - B. Such documentation should be sufficiently detailed that agencies can understand the underlying technical and analytical basis for the conclusions of the evaluations, testing, or red-teaming, and reproduce the results where appropriate, if practicable.
 - C. In addition to categories of risk that vendors include in evaluations, testing, and red-teaming, agencies should determine categories of risk covered in such documentation. Agencies are encouraged to require such documentation to cover evaluation, testing, or red-teaming performed on content falling under the following categories of risk:⁴⁰
 1. Inclusion of Chemical, Biological, Radiological, and Nuclear Information
 2. Confabulation (i.e., erroneous or false content)
 3. Dangerous or Violent Recommendations
 4. Data Privacy Risks
 5. Human-AI Configuration⁴¹
 6. Information Security Violations
 7. IP, including Copyright, Violations
 8. Obscene, Degrading, and/or Abusive Content, including CSAM and Non-Consensual Intimate Imagery
 9. Toxicity, Bias, and Homogenization⁴²
 - D. Agencies may accept documentation from the U.S. AI Safety Institute detailing the Institute’s testing of relevant AI systems as contributing to compliance with this subsection, though they must also require additional documentation as appropriate.
- vi. **Mitigate Environmental Impacts.** In accordance with OMB Memorandum M-24-10 Section 5(d)(viii), when procuring a general use enterprise-wide generative AI system or service, agencies should consider the degree of environmental impact of training and using the AI system or service. This can include requiring from vendors as part of a

⁴⁰ This list represents a subset of those identified in NIST’s AI RMF: Generative Artificial Intelligence Profile

⁴¹ This generally refers to interactions between humans and AI systems that can result in “algorithmic aversion, automation bias or over-reliance, misalignment or mis-specification of goals and/or desired outcomes, deceptive or obfuscating behaviors by AI systems based on programming or anticipated human validation, anthropomorphization, or emotional entanglement between humans and AI systems; or abuse, misuse, and unsafe repurposing by humans.” See NIST’s AI RMF Generative AI Profile.

⁴² Risks associated include those highlighted in NIST’s AI RMF Generative AI Profile, including “difficulty controlling public exposure to toxic or hate speech, disparaging or stereotyping content; reduced performance for certain sub-groups or languages other than English due to non-representative inputs; undesired homogeneity in data inputs and outputs resulting in degraded quality of outputs.”

solicitation and evaluation process, and through resulting contractual requirements where appropriate, documentation on the following to understand the nature of the generative AI system's environmental impact:

- A. Outlining and quantifying, to the best of the vendor's ability, the energy expended and projected for training and using the generative AI system. This could include considerations for carbon emissions and resource consumption, including energy and water usage, from supporting data centers. This could also include, when possible, considerations associated with the hardware necessary for the generative AI system, including raw mineral extraction and electronic waste generation.
- B. Demonstrating whether the vendor has implemented methods to improve the efficiency and sustainability of the generative AI system, such as the development and use of efficient data centers, the prioritization of renewable energy to supply data centers, partnerships with efficient data center providers (if the vendor is not a data center provider), and conducting research on how to make the system more efficient.
- C. Identifying opportunities for the generative AI system to be used for other agency or government needs in a manner that avoids the development of duplicative or redundant AI systems or services.

5. PROMOTING A COMPETITIVE AI MARKET WITH INNOVATIVE ACQUISITION

A diverse market of suppliers to enable Federal use of AI can strengthen agencies' ability to acquire the AI and supporting infrastructure they need, reduce vendor lock-in, help agencies lower costs, remove barriers for small businesses and potential new entrants, and build the economic strength of the United States.⁴³ Robust competition can help mitigate some risks associated with AI, as vendors compete to reduce risks to security and safety to improve the quality of their offerings. Because the AI marketplace is dynamic and evolving, as AI technology itself continues to evolve, agencies must, to the greatest extent practicable, ensure short-term and long-term interoperability across the elements of the technical stack⁴⁴ in a system. This section establishes requirements and recommendations around competition – with an emphasis on agency data management, portability, and related interoperability – and the use of innovative acquisition practices as a means of shaping and diversifying the AI marketplace, and ensuring agencies are able to acquire best-in-class AI, consistent with OMB Memorandum M-24-10. Appendix I details innovative practices and resources agencies are strongly encouraged to leverage.

a. Minimize Vendor Lock-In and Potential Switching Costs Through Contractual Requirements Language.

⁴³ OMB Memorandum, M-23-11, *Creating a More Diverse and Resilient Federal Marketplace through Increased Participation by New and Recent Entrants* (February 17, 2023), <https://www.whitehouse.gov/wp-content/uploads/2023/02/M-23-11-Creating-a-More-Diverse-and-Resilient-Federal-Marketplace.pdf>.

⁴⁴ The technical stack refers to the different layers at which the diverse tasks necessary to enable AI adoption occur. Supporting and performing these tasks includes data collectors and labelers, model developers, downstream developers, model hubs, hardware providers, compute providers, system integrators, and AI service providers.

Contractual requirements can help prevent vendor lock-in, which may arise from unreasonably high costs to switch vendors, insufficient knowledge transfer and documentation about acquired AI systems and services, unreasonable licensing restrictions, and maintenance and auditing practices intended to prevent agencies from switching vendors or reducing use of acquired AI and supporting infrastructure. Agency contractual requirements must, where relevant and practicable, reflect the following principles:

- i. Vendors commit to transfer knowledge to appropriate agency staff, including by demonstrating and providing information about the AI's operations, writing and committing to implementing a plan to train agency staff, and documenting information necessary for an agency to switch vendors or for a different vendor to build upon or integrate the AI system or elements within that system into a different system;
- ii. Vendors follow practices to promote data and model portability⁴⁵ (e.g., at the conclusion of a contract, vendors provide agencies with all of the relevant documentation as outlined in Section 4(c)(iii)(A-E) generated for the agency throughout the duration of the contract to ensure the agency's ability to continue operations);
- iii. Vendors provide agencies with appropriate rights to code, data, and models first produced in performance of the contract (see Section 4(c)(iii)),⁴⁶ addressing the Government's right to use, disclose, reproduce, prepare derivative works, distribute copies to the public, perform publicly and display publicly, in any manner and for any purpose, and to have or permit others to do so;
- iv. Vendors agree to transparent licensing terms that do not unduly restrict agencies' ability to use the acquired AI system or service, and to supporting infrastructure, in reasonably expected ways (see Section 4(c)(iii));
- v. Vendors agree to agency requests for pricing transparency across the total lifecycle of AI design, development, and duration of use that are domain and/or use case specific. This includes offering services without egress fees for cloud computing or other services that charge users for taking their data or business away from the service; and
- vi. Vendors do not employ contract terms that limit the sharing of the pricing information described in sub-section (a)(v) of this section with other agencies, except where sharing is prohibited by law, where the contract identifies the information as classified, or where the agency makes a determination approved by the agency Senior Procurement Executive (without delegation) after consultation with the Administrator for Federal Procurement Policy of a compelling business interest to restrict sharing.

b. Prioritize Vendor Practices that Increase Interoperability, Transparency, and Public Access.

⁴⁵ This generally refers to storing and representing data and models in a manner that allows for them to be re-used without an agency or other vendor having to perform data conversions or potentially building entirely separate redundant storage systems.

⁴⁶ This broadly included forms of data such as data from transactional and operational systems, metadata, and logging data first produced in performance of the contract

During market research, pre-solicitation, and evaluation, agencies should assess vendor approaches to interoperability, specifically as it pertains to data sharing, data portability, and model portability within different executive branch agency environments.⁴⁷ Emphasizing the strategic value of data access and the importance of data management⁴⁸ can help to minimize the risks of vendor lock-in, high switching costs in the event of switching vendors in future, or other negative competitive effects. Agencies are strongly encouraged to consider vendor practices that increase competition by including interoperability as part of evaluation criteria. Examples of relevant practices include, but are not limited to, the following:

- i. Well-defined application programming interfaces (APIs), particularly within acquired architectures, that promote interoperability with other elements of the technical stack;
- ii. Robust documentation regarding decisions related to foundational model development, coding languages used, testing scripts and protocols, and other decisions related to the development of AI tools in a developer experience framework that facilitates the transition of AI tools from one vendor to the next;
- iii. Open-source licenses to vendor's products, including AI models, AI systems, AI services, and datasets; and
- iv. Transparent and non-discriminatory pricing practices, including:
 - A. Offering products without bulk pricing arrangements, tying arrangements, steering arrangements, minimum spend requirements, or other agreements that encourage consolidation of spending with one vendor or one group of vendors through fixed contract lengths, exclusive discounts, or other incentives;
 - B. Offering systems or services at uniform and publicly available prices and not engaging in self-preferencing;
 - C. Providing equal access on comparable terms to downstream businesses, such as by refraining from self-preferencing vertically integrated systems or services; and
 - D. Providing information about which subcontractors, including system integrators, were engaged, how they were selected, and how their involvement impacts price.

c. Take Advantage of Innovative Acquisition Practices.

Leveraging innovative business practices and technologies to secure better contract outcomes is a guiding principle for the Federal Acquisition System.⁴⁹ The adoption and scaling of innovative techniques and technologies enables the acquisition workforce to execute agency missions more efficiently with increased customer satisfaction, better performance, and lower cost. Agencies must promote managed risk-taking and empower the acquisition workforce to test and use innovative acquisition approaches, with the necessary safeguards in place. This helps ensure agencies can evaluate vendors and proposals, get the best-suited AI for their needs, and reduce unnecessary burdens and duplicative development costs. Appendix I of this memorandum outlines practices and resources that agencies are strongly encouraged to leverage throughout the

⁴⁷ A-130

⁴⁸ A-137

acquisition lifecycle for AI solutions, consistent with the risk-management principles and practices described in the other sections of this document.

6. DEFINITIONS

The below definitions apply for the purposes of this memorandum:

Agency: The term “agency” has the meaning provided in 44 U.S.C. § 3502(1).

Artificial Intelligence (AI): The term “artificial intelligence” has the meaning provided in Section 238(g) of the John S. McCain National Defense Authorization Act for Fiscal Year 2019,⁵⁰ which states that “the term ‘artificial intelligence’ includes the following”:

1. Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.
2. An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.
3. An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
4. A set of techniques, including machine learning, that is designed to approximate a cognitive task.
5. An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.

For the purposes of this memorandum, the following technical context should guide interpretation of the definition above:

1. This definition of AI encompasses, but is not limited to, the AI technical subfields of machine learning (including deep learning as well as supervised, unsupervised, and semi-supervised approaches), reinforcement learning, transfer learning, and generative AI.
2. This definition of AI does not include robotic process automation or other systems whose behavior is defined only by human-defined rules or that learn solely by repeating an observed practice exactly as it was conducted.
3. For this definition, no system should be considered too simple to qualify as covered AI due to a lack of technical complexity (e.g., the smaller number of parameters in a model, the type of model, or the amount of data used for training purposes).
4. This definition includes systems that are fully autonomous, partially autonomous, and not autonomous, and it includes systems that operate both with and without human oversight.

AI Model: The term “AI model” has the meaning provided in Section 3(c) of Executive Order 14110.

⁵⁰ Pub. L. No. 115-232, § 238(g), <https://www.govinfo.gov/content/pkg/PLAW-115publ232/pdf/PLAW-115publ232.pdf>.

Artificial Intelligence System: The term “artificial intelligence system” has the definition provided in Section 7223 of the Advancing American AI Act, which states that “[t]he term ‘artificial intelligence system’—

(A) means any data system, software, application, tool, or utility that operates in whole or in part using dynamic or static machine learning algorithms or other forms of artificial intelligence, whether—

(i) the data system, software, application, tool, or utility is established primarily for the purpose of researching, developing, or implementing artificial intelligence technology; or

(ii) artificial intelligence capability is integrated into another system or agency business process, operational activity, or technology system; and

(B) does not include any common commercial product within which artificial intelligence is embedded, such as a word processor or map navigation system.”

Acquisition: The term “acquisition” has the meaning provided in the Federal Acquisition Regulation, subpart 2.1, which states that “‘acquisition’ means the acquiring by contract with appropriated funds of supplies or services (including construction) by and for the use of the Federal Government through purchase or lease, whether the supplies or services are already in existence or must be created, developed, demonstrated, and evaluated. Acquisition begins at the point when agency needs are established and includes the description of requirements to satisfy agency needs, solicitation and selection of sources, award of contracts, contract financing, contract performance, contract administration, and those technical and management functions directly related to the process of fulfilling agency needs by contract.” Procurement is defined in the Federal Acquisition Regulation, subpart 2.1, by reference to the definition of the term “acquisition.”

Confabulations: The term “confabulations” has the meaning provided in NIST AI 600-1, which states that the production of confidently stated but erroneous or false content (known colloquially as “hallucinations” or “fabrications”) by which users may be misled or deceived. Some commenters have noted that the terms “hallucination” and “fabrication” anthropomorphize generative AI, which itself is a risk related to generative AI systems as it can inappropriately attribute human characteristics to non-human entities.

Face recognition/face identification: The terms “face recognition” and “face identification” have the meaning provided in NIST’s ANSI/ASTM E2916-19e1 Standard Terminology for Digital and Multimedia Evidence Examination, which states that facial identification,(1) by automated systems, the automated searching of a facial image as a probe in a facial recognition system (one-to-many), typically resulting in a group (candidate list) of facial images being returned to a human operator in ranked order based on system-evaluated similarity; (2) by humans, the mental process by which an observer identifies a person as being one they have seen before.

Federal Information: The term “Federal information” has the meaning provided in OMB Circular A-130.

Generative AI: The term “generative AI” has the meaning provided in Section 3(p) of Executive Order 14110.

Hyperparameters: The term “hyperparameters” has the meaning provided in the NIST AI RMF glossary, which states that “hyperparameters” are “the parameters that are used to either configure a machine learning model (e.g., the penalty parameter C in a support vector machine, and the learning rate to train a neural network) or to specify the algorithm used to minimize the loss function (e.g., the activation function and optimizer types in a neural network, and the kernel type in a support vector machine).”

National Security System: The term “National Security System” has the meaning provided in 44 U.S.C. § 3552(b)(6).

Red-teaming: The term “red-teaming” has the meaning provided in Section 3(d) of Executive Order 14110.

Rights-Impacting AI: The term “rights-impacting AI” has the meaning provided in OMB Memorandum M-24-10.

Risks from the Use of AI: The term “risks from the use of AI” has the meaning provided in OMB Memorandum M-24-10.

Safety-Impacting AI: The term “safety-impacting AI” has the meaning provided in OMB Memorandum M-24-10.

Appendix I: Promoting Innovative Acquisition

This appendix builds upon actions taken by OMB and the CAO Council to create an innovation-friendly acquisition environment that empowers, enables, and encourages members of the acquisition workforce to test and share new and better ways of conducting acquisitions with the help of acquisition innovation advocates at every cabinet-level department and agency, and a growing number of labs, coaches, and workforce development tools and programs.⁵¹

It outlines actions agencies should take to promote innovative practices that can lower barriers to entry for small businesses and other vendors, and share information on their use of AI-driven solutions to promote responsible adoption and adaption without redundant investment. Agencies are also encouraged to consider how their acquisitions can be used to promote competitive markets including through considerations of the extent vendors are already dominant players, are highly vertically integrated, limit interoperability and data portability, or engage in self-preferencing. This appendix also describes actions to ensure the acquisition workforce has access to the resources and the training it needs to successfully apply innovative techniques consistent with applicable guardrails described in this memorandum. Agencies are strongly encouraged to adopt the following practices:

a. Take Advantage of Performance-Based Acquisition Techniques.

Agencies are strongly encouraged to use performance-based requirements that allow agencies to understand and evaluate vendor claims about their AI systems or services prior to contract award, acquire AI capabilities that address their needs, and better perform post-award monitoring. Performance based techniques include:

- i. Statements of Objectives (SOOs) and Performance Work Statements (PWS), which provide agencies with more flexibility to acquire AI systems or services that meet agencies' outcome-oriented needs but may not meet unnecessary or overly-limiting requirements in Statements of Work (SOW).
- ii. Quality Assurance Surveillance Plans (QASP), which can help agencies overcome challenges in defining relevant performance metrics pre-solicitation and can enable a more collaborative process for negotiating a QASP that meets agency needs and objectives. Government personnel should be prepared to assume a more active role in performance monitoring.
- iii. Contract incentives to improve the performance for AI systems or services and their interoperability such that other agencies and vendors can improve upon, replicate, and further develop existing federal AI source code. Incentives can be based on metrics and provisions in QASPs. When determining whether to include performance-based

⁵¹ For a comprehensive overview of steps that have been taken to promote innovation in acquisition, see [Acquisition Innovation & Small Business Participation in Federal Procurement](#). Efforts to promote an innovative mindset that values creative thinking, outcomes, and risk management over compliance are contributing to a greater sense of empowerment to problem solve within the acquisition workforce. According to Federal Employee Viewpoint Surveys (FEVS) conducted between 2017-2022, the acquisition workforce has consistently responded more favorably than the workforce at large to questions about feeling empowered to use good business judgment in meeting the daily responsibilities of government service

incentives, agencies must carefully consider whether the established metrics are correctly tied to desired business and societal outcomes, and whether they can adequately measure baseline performance of the AI systems or services.

b. Enlist Acquisition Coaches.

Agencies are encouraged to include an acquisition innovation coach, or an individual with similar expertise, as a consultant to acquisition teams, such as Integrated Program Teams (IPTs)⁵² or the source selection evaluation board. Acquisition coaches are individuals identified by the agency's acquisition innovation advocate who have developed expertise in using innovative acquisition techniques such as those on the Periodic Table of Acquisition Innovation (PTAI),⁵³ and who can enable efficiencies in the acquisition process. Coaches can provide instruction on the application of strategies through consultation, hands-on testing of practices, and related assistance.⁵⁴

c. Perform Thorough Market Research.

Prior to acquiring AI, agencies should regularly leverage collective knowledge-sharing platforms and portals across the Federal Government, such as the AI acquisition resources housed on GSA's platform identified in Section 3(b)(i) of this memorandum. Market research strategies identified on the PTAI, such as industry one-on-ones, reverse industry days, demonstrations, and interactive Q&A, help the government and the marketplace exchange information upfront. Agencies are encouraged to consider market research strategies geared towards identifying small business new entrants. Such techniques to conduct this market research include:

- i. Reviewing AI systems or services pre-screened by the Department of Defense on its Tradewinds Solutions Marketplace;
- ii. Researching AI systems or services offered by small businesses in the Small Business Innovation Research (SBIR) program. While agencies are encouraged to consider efforts, as appropriate, in all three phases of the SBIR program, businesses in Phase III of the program have solutions with previously demonstrated technical merit, feasibility, and commercial potential that may be acquired by any agency; and

⁵² IPTs are designed to include appropriate cross-functional expertise (e.g., acquisition, IT, cybersecurity, privacy, civil rights and civil liberties, budgeting, data, legal, program evaluation). For more information on Specialized IT Acquisition Cadres and Integrated Product Teams, see generally, Memorandum from the Admin. for Fed. Procurement Policy & U.S. Chief Inf. Officer, *Acquisition Innovation Labs & Pilot for Digital Acquisition Innovation Lab* (Mar. 9, 2016), <https://techfarhub.usds.gov/assets/files/AIA-March9Memo.pdf>, and GAO-17-8, *IT WORKFORCE: Key Practices Help Ensure Strong Integrated Program Teams; Selected Departments Need to Assess Skill Gaps* (Nov. 2016), <https://www.gao.gov/assets/d178.pdf>.

⁵³ The PTAI is the governmentwide knowledge management portal for acquisition innovation. See <https://acquisitiongateway.gov/periodic-table>.

⁵⁴ For additional information on developments in the use of acquisition innovation coaches, see section II.C.3. of Acquisition Innovation & Small Business Participation in Federal Procurement.

- iii. Pursuing opportunities to leverage third party experts to research and connect agencies with contractors that are suited to their particular needs through the use of Partnership Intermediary Agreements, or similar types of agreements.

d. Conduct Informed Analyses of Vendor Proposals.

To assess vendor capabilities and evaluate whether a vendor’s proposed solution is suited for agency needs, agencies should use “show, don’t tell” approaches that require vendors to demonstrate capabilities during solicitation, evaluation, and source selection. Example approaches identified on the PTAI include:

- i. Oral presentations, which allow the Government to hear directly from the vendor and their technical experts about their solutions, and technical demonstrations, that give government evaluators the opportunity to see the proposed solution in action; and
- ii. Pilots, trials, and other proofs of concept involving multiple vendors that promote diversity of approaches to foster competition. By using pilots to help the government learn about and test a capability in performing to specific use cases, based on performance-based requirements, agencies can better ensure over time that the acquired AI is best equipped to deliver the service needed.

e. Leverage an AI-Trained Workforce.

Agencies should ensure not only that their acquisition workforce participates in AI-related training, but also that trained employees are able to effectively apply knowledge from those trainings to the acquisition of AI systems or services. Agencies should incentivize employees to participate in training, including training established and provided under the AI Training for the Acquisition Workforce Act⁵⁵ (AI Training Act), to enable their participation in IPTs, implement novel acquisition innovations and approaches listed in Section 4(a)(ii) of this memorandum. As such, training should:

- i. Cover topics such as identifying AI features or a vendor’s use of AI; evaluating vendor claims; scoping, acquiring, testing and de-biasing AI systems or services;
- ii. Address criteria for how the agencies can identify limitations of interoperability and data portability;
- iii. Emphasize the importance of open-source best practices as a method to promote interoperability, mitigate vendor lock-in, and enable other agencies and their vendors to build upon existing code and leverage lessons learned through testing and development efforts;
- iv. Include instructions for designing and developing the integration of AI in a way that considers impact to Federal workers; and

⁵⁵ AI Training Act, Pub. L. No. 117-207 (2022).

- v. Create necessary resources for Federal workers to learn about and provide feedback on how the AI integration impacts their workflow.

f. Use Proven Strategies to Promote Competition.

The following list highlights practices that are designed to help agencies leverage the power of competition to obtain better value and results from their AI acquisitions:

- i. Provide performance-based criteria in contract language as requirements for continued use;
- ii. Limit the duration of contracts to ensure regular re-assessment in alignment with fulfilling the agency's needs;
- iii. Carefully separate and intentionally diversify the different technical elements that compose an AI system or service (e.g., data labelers, system integrators) through multiple contracts and/or modular/multi-award contracts or teaming agreements;
- iv. Describe requirements in vendor- and model-agnostic ways, enabling projects to quickly test and switch models and capture the benefits of advances by providers, including cost, quality, speed, new functionality, or switching to an agency-trained and -hosted model in the future;
- v. Utilize prototypes, challenge-based prize competitions,⁵⁶ tech sprints, pilot programs, and industry engagement days to test offeror capabilities prior to purchase to prioritize performance metrics in data portability and attract non-incumbent vendors in evaluation process;
- vi. Address in acquisition plans for new multiple-award contracts the use of on-ramps, which allow for small and large businesses to be added during the performance period for long-term contracts, which under FAR 19.301-2(a) are those contracts that are more than five years in duration, including options;⁵⁷
- vii. Prioritize the technical capability of vendors rather than defaulting to lowest price technically acceptable analysis, such as by including consideration for interoperability toward future or related use cases;
- viii. Explore authorities available to the agency in addition to those provided by the FAR, such as commercial services opening pilots and other transactions authorities, when they may be more suitable than the FAR in lowering barriers to entry;

⁵⁶ Ali Crawford and Ido Wulkan, "Federal Prize Competitions" (Center for Security and Emerging Technology, November 2021), <https://doi.org/10.51593/2021CA002>

⁵⁷ Memorandum for Chief Acquisition Officers and Senior Procurement Executives, Increasing Small Business Participation on Multiple-Award Contracts (January 25, 2024), https://www.whitehouse.gov/wp-content/uploads/2024/01/REV_Increasing-Opportunities-to-Small-Businesses-under-MACs-CATS-Final-Copy-1-25-24.pdf. For examples of on-ramps in the context of multiple-award contracts, see "Use Cases and Documentations" under the "On/Off Ramp" tile of the Periodic Table of Acquisition Innovations.

- ix. For agencies that are able, utilize SBIR grants to fund pilots that incorporate agile new entrants;
- x. Explore the offerings of non-incumbent providers of AI systems or services, even when current providers offer similar features;
- xi. Explore the use of on-premises and localized cloud environments (such as those provided through the National AI Research Resource Pilot program or the National Secure Data Service Demonstration Project) to test AI models and evaluate vendor claims before acquisition, so as to develop a standardized adopt, trial and assess framework for AI acquisition;
- xii. Explore the use of building dedicated, shareable hardware environments for compute (such as procuring Graphics Processing Units, Tensor Processing Units, etc.) to build capabilities of AI testing and development that do not rely on API access through existing cloud service providers, especially for testing and implementing generative AI systems; and
- xiii. Safeguard against anticompetitive practices such as steering arrangements, tying arrangements or other forms of self-preferencing across services, including by ensuring a competitive process to avoid entrenching dominant players. As with procuring any service, agencies should also protect against collusion, such as market allocation, bid rigging, and price fixing. Suspected collusion can be reported to the DOJ Antitrust Division's [Procurement Collusion Strike Force](#).

Appendix II: Implementation Schedule for Requirements

Action	Section	Deadline
Agencies shall identify any contracts associated with agency use of rights-impacting or safety-impacting AI.	2(c)(i)	November 1, 2024
Agencies shall ensure that contracts identified as associated with agency use of rights-impacting or safety-impacting AI are brought into compliance with the requirements of M-24-10 and Sections 4(d), 4(e), and 4(f)(ii) of this memorandum.	2(c)(ii)(A) and 4(d), 4(e), 4(f)(ii)	December 1, 2024
Agencies shall ensure that new contracts issued in support of agency use of rights-impacting or safety-impacting AI are consistent with the requirements of M-24-10 and this memorandum.	2(c)(ii)(B)	December 1, 2024
Each agency must establish or update policies and procedures for internal agency collaboration to ensure that acquisition of an AI system or service will have the appropriate controls in place to comply with the requirements of this memorandum and that the agency's use of the acquired AI will conform to OMB Memorandum M-24-10.	3(a)(i)	Ongoing
Agency CAIOs must submit written notification to OMB identifying progress made toward implementing cross-functional collaboration, and identifying any challenges or best practices.	3(a)(i)	March 23, 2025
Each agency CAIO must submit to OMB a plan for ensuring that the CAIO coordinates on AI acquisition with the CAO, CIO, CISO, CFO, SAOP, and other relevant officials (e.g., Chief Technology Officer, Chief Data Officer, Chief Competition Officer, an official responsible for protecting civil liberties).	3(a)(ii)	March 23, 2025
Agencies shall ensure that the SAOP and agency privacy program are involved during pre-solicitation acquisition planning and while the agency is defining the requirements for prospective AI systems or services.	4(b)(i)(A)	Ongoing
Agencies shall equip their privacy programs with the necessary resources to develop and implement, in coordination with other agency officials responsible for AI acquisition, appropriate contractual terms and conditions, policies, and processes to ensure that agency AI contracts comply with the privacy requirements in law and policy when they involve creating, collecting, using, processing, storing, maintaining, disseminating, disclosing, or disposing of	4(b)(i)(B)	Ongoing

information on behalf of a Federal agency or operating or using information systems on behalf of a Federal agency.		
Agencies should ensure that contractual requirements address risks inherent in the procurement of AI systems that identify individuals using biometric identifiers.	4(b)(ii)	Ongoing
Agencies must ensure that performance-based requirements identified by the agency ensure, to the greatest extent practicable, that the AI system or service to be acquired will be appropriate for the expected use contexts.	4(c)(ii)	Ongoing
Agencies must include appropriate contractual terms that clearly delineate the respective ownership and IP rights of the Government and the contractor.	4(c)(iii)	Ongoing
Agencies must develop an approach to IP that considers what rights and deliverables are necessary for the agency to successfully accomplish its mission, protects Federal information used by vendors in the development and operation of AI systems and services for the Federal Government, promotes the development of AI code, avoids vendor lock-in, and avoids unnecessary costs.	4(c)(iii)	Ongoing
Agencies must scrutinize terms of service and licensing terms to ensure that they clearly articulate the scope of rights needed by the Government over its own data and any derived products.	4(c)(iii)	Ongoing
Agencies must ensure that terms of contracts explicitly address how a vendor ensures (e.g., through a quality management system) compliance with relevant data management directives and policies.	4(c)(iv)	Ongoing
Agency contractual requirements must, where relevant and practicable, reflect the principles outlined in Section 5(a) of this memorandum.	5(a)	Ongoing
Agencies must promote managed risk-taking and empower the acquisition workforce to test and use innovative acquisition approaches, with the necessary safeguards in place.	5(c)	Ongoing

Additional requirements for rights-impacting and safety-impacting AI

Action	Section	Deadline
Where practicable, agencies must disclose in solicitations whether the planned use is rights-impacting or safety-impacting.	4(d)(i)	Ongoing

Agencies must ensure that vendors provide them with the information and documentation necessary to monitor the performance of an AI system or service and implement applicable requirements of OMB Memorandum M-24-10, as delineated in subsections 4(d)(ii)(A-H).	4(d)(ii)	Ongoing
Agencies must ensure that contractual terms provide the ability to regularly monitor and evaluate (e.g., on a quarterly or bi-annual basis, based on the needs of the program) performance and risks throughout the duration of the contract, as delineated in subsections 4(d)(iii)(A-F).	4(d)(iii)	Ongoing
Agencies must have the ability, throughout the entire lifecycle of the contract, to update risk mitigation options and prioritize performance improvement of the AI system or service.	4(d)(iv)	Ongoing
Agencies must, to the greatest extent practicable, include contractual terms requiring that vendors have a process for identifying and disclosing to agencies serious AI incidents and malfunctions of the acquired AI system or service.	4(d)(v)	Ongoing

Additional requirements for rights-impacting AI

Action	Section	Deadline
To the extent practicable and appropriate, agencies must identify in requirements documents for rights-impacting AI systems or services whether use of the AI will involve notifying individuals of AI-enabled decisions affecting them and affording opportunities for human consideration and remedy.	4(e)(i)	Ongoing
Agencies must include in contracts any requirements for additional access, information, or necessary documentation about the AI system or service necessary for the agency to carry out any plans for notice and appeal procedures.	4(e)(ii)	Ongoing

Additional requirements for general use enterprise-wide generative AI

Action	Section	Deadline
Agencies must include contractual terms which ensure that any audio, image, and video outputs of AI systems are not readily distinguishable from reality are created or modified using mechanisms that allow the outputs to be identified as generated by AI, attributed to the specific model that was used to produce the	4(f)(i)(A)	Ongoing

output, and linked with other relevant information about the origin or history of outputs.		
Agencies must document how the general use enterprise-wide generative AI was or will be trained and evaluated, including relevant information about data, data labor, compute, model architecture, and relevant evaluations.	4(f)(i)(B)	Ongoing
Agencies must consider including contractual requirements that ensure vendors provide appropriate protections, where practicable, against the AI systems or services being used in ways that are contrary to law and policy.	4(f)(ii)	Ongoing
Agencies must include contractual requirements which ensure that vendors employ best practices to prevent the AI system from generating toxic, false, illegal, violent, or otherwise harmful content; make best efforts to filter out known CSAM and NCII in their training data; and provide the agency with documentation on the components and capabilities in the model or system put in place by the vendor or configurable by agencies to limit harmful and illegal outputs.	4(f)(iii)	Ongoing
Agencies must take appropriate steps to provide an objective and empirical basis for competitively selecting the solution or solutions that provide the most value to the agency.	4(f)(iv)	Ongoing
Agencies must include contractual terms requiring vendors provide documentation on evaluations, testing, and red-teaming as delineated in subsections 4(f)(v)(A-D).	4(f)(v)	Ongoing