



Comments on Proposed OMB Memorandum, “Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence” OMB-2023-0020-0001

December 5, 2023

Workday appreciates the opportunity to provide the Office of Management and Budget (OMB) comments on the recently issued Proposed Memorandum, “Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence” (Draft Memorandum). When finalized, the Draft Memorandum will play an important role in how federal agencies can responsibly use and benefit from these critical emerging technologies. As written, the Draft Memorandum strikes a fair balance between advancing responsible AI and the federal government’s longstanding efforts to modernize its information technology systems.

We offer the following recommendations to strengthen OMB’s Draft Memorandum as it seeks to advance both of these goals:

- **Emphasize the importance of “buy, not build.”** Commercial solutions, including Workday’s, are better positioned than custom-built ones at helping agencies streamline and standardize administrative processes, gain unprecedented visibility across organizations, and adapt quickly to change, including changes in fast-evolving technologies such as AI.
- **Retain the current focus on the federal workforce.** OMB rightly identifies workforce challenges as a barrier to AI adoption. The federal government’s emerging focus on skills-based hiring, rather than solely on traditional credentials, is the right lens through which to view these workforce priorities. A skills-based approach to talent is difficult to scale without AI tools implemented in a responsible manner.
- **Impact assessments are the right AI accountability tool.** We applaud OMB’s decision to require federal agencies to conduct impact assessments on high-risk AI tools. We recommend focusing impact assessments on identifying, documenting, and mitigating the “reasonably foreseeable” risk of harm to an individual’s rights or safety. To streamline the impact assessment process, we recommend the use of common forms across federal agencies, which will rely on information regarding the quality and appropriateness of data in a standardized fashion.
- **Require federal agencies to use the NIST AI Risk Management Framework.** The framework was developed at the direction of Congress through an open, multi-stakeholder process with participation from academia, civil society, and the business community. There is growing bipartisan support in Congress for its adoption by the federal government. We urge OMB to go further than recognizing the NIST AI Risk Management as a “best practice” and require its adoption.
- **Refine the definition of “rights-impacting” AI.** We commend OMB for adopting a risk-based approach to AI governance. At the same time, we suggest that it refine the definition of “rights-impacting” AI by: (1) adopting an exhaustive and clearly defined list of

consequential *decisions*; and (2) focusing on AI tools that *automate* these decisions. Doing so would enable federal agencies to leverage human-in-the-loop AI systems and focus their risk management efforts on the tools with the highest risk profiles.

- **Distinguish between mission-specific and enterprise AI systems.** Federal agencies are best equipped to determine if an AI tool is “rights-” or “safety-impacting,” when it relates to their mission. By contrast, enterprise AI systems, such as those for human resources and financial management, would benefit from a uniform determination across the federal government. This will promote consistency in determinations of “rights-” and “safety-impacting” AI tools and encourage a more consistent level of protection.

Introduction & Background

About Workday

By way of introduction, Workday is a leading provider of enterprise cloud applications for finance and human resources, helping customers adapt and thrive in a changing world. Our applications for financial management, human resources, planning, spend management, and analytics are built with artificial intelligence (AI) and machine learning at the core to help organizations embrace the future of work. Workday is used by more than 10,000 organizations around the world and across industries—from medium-sized businesses to more than 50 percent of the *Fortune* 500. Our customers also include state, local, and county governments and institutions of higher education. The Workday customer community has 65 million users, and in April of this year, nearly one in four of all U.S. job openings was processed on the Workday platform.¹

Workday & the Federal Marketplace

In 2020, Workday officially entered the U.S. federal marketplace with our first federal agency customer, and in 2022, we obtained our FedRAMP authorization.^{2 3} Our entry comes at a time when the White House, Congress, and federal agencies are grappling with significant challenges, including recruiting, attracting, and retaining talent, reducing the government’s dependency on outdated, legacy technology, and bolstering cybersecurity. These growing challenges have occurred against the backdrop of an aging incumbent workforce and difficulty in attracting younger talent. By the end of 2023, nearly two-thirds of people in senior government leadership positions will reach retirement age. By contrast, employees who are millennials and Generation Z are only 7 percent of the federal workforce.⁴ This is compared to nearly 25 percent in the private sector.⁵

¹ Landman, Inna. “First Half 2023 Hiring Trends: Slowdown and Stagnation,” Workday Blog, September 20, 2023, <https://blog.workday.com/en-us/2023/first-half-2023-hiring-trends-slowdown-stagnation.html>.

² “Workday Achieves Fedramp Authorized Designation.” Workday Newsroom. <https://newsroom.workday.com/2022-07-13-Workday-Achieves-FedRAMP-Authorized-Designation>.

³ Robinson, Doug. “Tipping Point - Modernizing the Federal Workforce.” POLITICO, July 11, 2022. <https://www.politico.com/sponsor-content/2022/07/11/tipping-point-modernizing-the-federal-workforce>.

⁴ “Federal Workforce Needs to Be a Priority — No Really, a Top Priority This Time, Advocates Say.” 2020. Federal News Network. November 9, 2020.

<https://federalnewsnetwork.com/federal-report/2020/11/federal-workforce-needs-to-be-a-priority-no-really-a-top-priority-this-time-advocates-say/>.

⁵ Mace, Chris. “Budget Underscores Lack of Younger Federal Employees.” FEDweek. February 26, 2020. <https://www.fedweek.com/fedweek/budget-underscores-lack-of-younger-federal-employees/>.

Federal hiring managers often lack the data, reporting, and analytic tools to understand their workforce, and track and match skills so that federal employees can continue to grow their careers within the government.⁶ The federal government also often lacks systems to assess retention risks, provide development opportunities, and support self-service. Missing these essential tools can impact motivation, morale and engagement, even in a mission-driven environment.

Federal agencies realize they need to make investments in hiring, onboarding, and talent development initiatives and to create an environment more conducive to learning, growth, and equity.⁷ They are also looking to establish thoughtful leadership practices around diversity and inclusivity. As remote work continues and evolves, federal agencies will also be looking for better ways to help assess employee engagement, productivity, and well-being while maintaining a supportive work environment.

Given the need to focus on their primary missions, agencies can and should look to commercially available solutions to their enterprise technology needs, rather than build their own. There is no doubt that technologies developed by the private sector evolve and improve faster than in-house and custom-agency projects. This is particularly true for AI, whether integrated into enterprise software services or offered as a stand-alone tool. Increasing partnerships between the public and private sectors will reduce complexity and costs, allowing the federal agencies to focus resources and staff on their core missions.

Modern and secure cloud-based systems are integral to achieving these government-wide priorities. Workday's solutions unify financial and talent management, time and attendance, and recruitment, which allows agencies the ability to automate their day-to-day personnel processes and give their workers more control through self-service options.⁸ Our solutions enable agencies to operate more strategically, hire and develop a top-performing workforce and adapt quickly to change.

Skills to Empower the Federal Workforce

As with earlier advances in technology, AI will impact how people work and the skills their jobs require. Notable developments around generative AI are also accelerating the pace and depth of transformation that we will likely see in the next two to three years. As Workday partners with federal agencies to solve a litany of challenges, including strengthening and empowering the federal workforce, we recommend a shift to a skills-based approach to talent to adapt to the changing workplace.^{9 10 11}

⁶ "The Federal Government Must Transform to Attract Tomorrow's Workers." Workday Blog, October 18, 2023. <https://blog.workday.com/en-us/2023/government-transforms-to-attract-workers.html>.

⁷ Strengthening the federal workforce - The White House, n.d. https://www.whitehouse.gov/wp-content/uploads/2022/03/ap_7_strengthening_fy2023.pdf.

⁸ "Government Cloud Computing and Solutions." Workday. <https://www.workday.com/en-us/solutions/industries/public-sector/workday-government-cloud.html>.

⁹ Schlapp, Pete. "A.I. Is a Must for Skills-Based Organizations That Want to Move at the Speed of Future Business," Fortune, February 17, 2023, <https://fortune.com/2023/02/17/workday-future-business-skills/>.

¹⁰ Somers, David. "How Workday Is Delivering Next-Generation Skills Technology at Scale," Workday Blog, September 28, 2022, <https://blog.workday.com/en-us/2022/how-workday-delivering-next-generation-skills-technology-scale.html>.

¹¹ Bryan Hancock et al., "Taking a Skills-Based Approach to Building the Future Workforce," McKinsey & Company, November 15, 2022,

By a skills-based approach to talent, we mean an emphasis on what a person can do or learn, rather than solely on their credentials. Awareness of the benefits of a skills-based approach to hiring, learning, and career development is growing.¹² In hiring guidance issued by the Office of Personnel Management last year, the federal government recognized the importance of a skills-based approach to recruiting its workforce.¹³ Governors of at least ten states are also taking steps to remove degree requirements for most state opportunities.¹⁴

At Workday, we've found that successfully implementing a skills-based approach to talent can be difficult to scale without the right technology.¹⁵ People use different words to describe a skill, and very different skills can be described with the same word. Workers often struggle to identify which skills they should develop to advance their careers, and the lack of consistency makes it difficult for employers to identify workers who can fill an open role.

Fortunately, AI is unique in that it will drive change in the workplace and power the tools that workers and employers need to successfully navigate those changes. AI can process large amounts of data associated with occupational roles and responsibilities and develop so-called "ontologies" or vocabularies that make skills data actionable. Workday's Skills Cloud, for example, aligns skills to a common vocabulary by using machine learning to map how different skills relate to each other and evolve over time.¹⁶ Skills Cloud has been used over 40 million times, including by hiring managers for new job postings and by incumbent workers and candidates to communicate the skills they have. Over 25 percent of *Fortune* 500 companies are now live on Skills Cloud and workers have entered over 200 million skills into their profiles.

Incumbent workers can also use Career Hub, a one-stop-shop on the Workday platform where employees can find AI-enabled personalized recommendations, such as learning content and short-term projects on other teams where they can pick up new skills. For example, a program manager at the Department of Energy interested in management can discover a leadership role in another program office and take suggested learning courses to prepare them for that opportunity.¹⁷ The result is a win-win for workers and employers: the program manager can pursue growth opportunities that align with their career goals, and the federal government can benefit from and support their incumbent workforce. This is one example of how AI and skills can take the guesswork out of workforce development and facilitate data-driven reskilling.

Our Commitment to Responsible AI

<https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/taking-a-skills-based-approach-to-building-the-future-workforce>.

¹² Ferguson, Stephanie. "Understanding America's Labor Shortage," U.S. Chamber of Commerce, October 23, 2023, <https://www.uschamber.com/workforce/understanding-americas-labor-shortage>.

¹³ "OPM Releases Skills Based Hiring Guidance." U.S. Office of Personnel Management. <https://www.opm.gov/news/releases/2022/05/release-opm-releases-skills-based-hiring-guidance/>.

¹⁴ "States Consider Elimination of Degree Requirements." National Conference of State Legislatures. <https://www.ncsl.org/education/states-consider-elimination-of-degree-requirements>.

¹⁵ "Skills credentials and Workforce of the Future." Workday.

<https://www.workday.com/content/dam/web/en-us/documents/whitepapers/skills-credentials-and-workforce-of-the-future.pdf>.

¹⁶ Stratton, Jim, David Somers, Rowan Miranda, et al. "The Foundation of the Workday Skills Cloud." Workday Blog, 2020. <https://blog.workday.com/en-us/2020/foundation-workday-skills-cloud.html>.

¹⁷ Ernst, Chris. "Making the Shift to a Skills-First People Strategy." SHRM, August 23, 2022.

<https://www.shrm.org/executive/resources/people-strategy-journal/summer2022/pages/chris-ernst-workday-skills-first-people-strategy.aspx>.

At Workday, we see incredible opportunities for AI to unlock human potential, including in the federal enterprise. We also recognize that the risk of unintended consequences stemming from AI's use is real and that trust is foundational to any technology's adoption. As a cloud-native enterprise software company, we learned early on that investing in rigorous technology governance aligns our solutions with our company values and earns and retains our customers' trust.^{18 19} This is why Workday put in place a robust responsible AI program that includes:

- **Leadership commitment** from a Responsible AI Advisory Board that is led by our General Counsel and counts our Chief Compliance Officer, Chief Technology Officer, and Chief Diversity Officer among its members.
- **An independent Responsible AI team** of social and data scientists and technology experts that reports through our Chief Compliance Officer to our Board of Directors. The team develops and maintains Workday's AI governance framework and receives cross-company support, including from responsible AI "champions" who provide subject matter expertise so that AI products are developed in accordance with Workday's ethics principles.²⁰
- **Responsible AI guidelines and review processes** that turn our principles into documented practices and assessments. Our product teams use these assessments to evaluate a potential AI feature's risk profile in the earliest stages of its development. AI tools that may be used for consequential decisions, such as hiring or promotion, are treated as high-risk.
- **Disclosure** to equip our customers with a clear understanding of how our AI tools are developed and assessed. We also provide our customers with disclosures and choices about how their data is used.

Workday has instituted these guardrails so that our AI tools are developed in a responsible manner. In addition to disclosure, we provide our customers with the means to access their own data for bias testing and the choice of whether to use an AI tool at all. In this way, we view AI governance as a partnership between *developers*, or organizations that produce or design AI tools, and *deployers*, which use AI tools and interact with end users.²¹ Similar to privacy and data protection laws, AI governance frameworks should distinguish between these roles and assign obligations accordingly. In the context of AI risk management, the Organization for Economic Cooperation and Development (OECD) has recognized this distinction as AI evaluated "in the lab" (i.e., developers) versus AI evaluated "in the field" (i.e., deployers).²² Understanding the differences in control and visibility between developers and deployers is essential for a workable life-cycle approach to AI risk management.

Workday understands that the lack of public trust in AI must be addressed across industry and in partnership with governments. We know that people and organizations will not use technology

¹⁸ "Our Core Values." Workday. <https://www.workday.com/en-us/company/about-workday/core-values.html>.

¹⁹ Cosgrove, Barbara. "Safeguarding Privacy while Innovating with AI." Workday Blog, May 24, 2023. <https://blog.workday.com/en-us/2023/safeguarding-privacy-while-innovating-ai-workday.html>.

²⁰ Trindel, Kelly. "Workday's Continued Diligence to Ethical AI and ML Trust." Workday Blog, September 19, 2023. <https://blog.workday.com/en-us/2022/workdays-continued-diligence-ethical-ai-and-ml-trust.html>.

²¹ "AI Developers and Deployers - An Important Distinction." BSA, March 26, 2023. <https://www.bsa.org/policy-filings/ai-developers-and-deployers-an-important-distinction>.

²² "OECD Framework for the Classification of AI Systems." OECD, February 22, 2022. <https://www.oecd.org/publications/oecd-framework-for-the-classification-of-ai-systems-cb6d9eca-en.htm>.

they don't trust. *Workday strongly supports government action to cultivate trust in AI, including new regulations on AI developed and used for consequential decisions*, such as decisions to hire, promote, or terminate an employee. We were early champions for the creation of the National Institute of Standards and Technology's (NIST) AI Risk Management Framework (AI RMF or Framework) because we recognized the need for a robust, flexible, and commonly accepted benchmark for AI governance.²³ Workday is an early adopter, and in September, NIST published a case study featuring Workday's use of the AI RMF, the first of any organization.²⁴

As policymakers began considering a path forward on AI policy, we recognized the need for a practical roadmap for responsible AI in the workplace. With this in mind, Workday joined the Future of Privacy Forum and other leading technology companies to co-develop the *Best Practices for AI and Workplace Assessment Technologies*.²⁵ The *Best Practices* leverage our experience delivering trustworthy enterprise AI capabilities at scale, as well as federal government benchmarks, including the NIST AI RMF and guidance from the Equal Employment Opportunity Commission and Federal Trade Commission (FTC).²⁶ We were proud to endorse the *Best Practices* and have called on other organizations to join us in using them.²⁷

Specific Comments on the Draft Memorandum

Advancing Responsible AI Innovation

Workday applauds OMB for recognizing that "AI can improve operations and deliver efficiencies across the federal government." We strongly agree with the Draft Memorandum's focus on identifying and removing barriers to the responsible adoption of AI in the federal enterprise. Robust AI governance is *not* incompatible with the federal government's longstanding efforts to modernize its information technology systems. The Draft Memorandum strikes a fair balance between these policy aims. We offer the following recommendations with the goal of advancing both priorities.

OMB should emphasize the importance of "buy, not build."

²³ "A Timely Bipartisan Push for Trust in AI: Congress and the NIST Trustworthy AI Framework." Morning Consult, January 12, 2021.

<https://morningconsult.com/opinions/congress-and-the-nist-trustworthy-ai-framework/>.

²⁴ "Using the AI Risk Management Framework - Workday." NIST, 2023.

<https://www.nist.gov/system/files/documents/2023/09/14/workday-success-story-final-for-release.pdf>.

²⁵ "Future of Privacy Forum and Leading Companies Release Best Practices for AI in Employment Relationships." Future of Privacy Forum, September 19, 2023.

<https://fpf.org/blog/future-of-privacy-forum-and-leading-companies-release-best-practices-for-ai-in-employment-relationships/>.

²⁶ "Assessing Adverse Impact in Software, Algorithms, and Artificial Intelligence Used in Employment Selection Procedures under Title VII of the Civil Rights Act of 1964." US EEOC.

https://www.eeoc.gov/laws/guidance/select-issues-assessing-adverse-impact-software-algorithms-and-artificial?utm_content=&utm_medium=email&utm_name=&utm_source=govdelivery&utm_term=.

²⁷ "Future of privacy forum and leading companies release best practices for AI in employment relationships." Future of Privacy Forum.

<https://fpf.org/blog/future-of-privacy-forum-and-leading-companies-release-best-practices-for-ai-in-employment-relationships/>.

In 2022, there was \$47.4 billion of private investment made in AI in the United States.²⁸ This underscores the scale of resources the U.S. private sector is devoting to developing cutting-edge AI technologies and deploying them at scale. For the federal government to efficiently and effectively modernize its IT systems, it is critical that federal agencies leverage the commercial solutions that emerge from these investments.

With this in mind, Workday welcomes the Draft Memorandum's recognition in Section 4 that "AI can improve operations and deliver efficiencies across the Federal Government." However, the Draft Memorandum also asserts that agencies "should *build* internal enterprise capacity to support responsible AI innovation." Unfortunately, this language may be interpreted by agencies as a directive to overlook commercially available solutions in favor of developing and deploying their own AI tools. *We urge OMB to clarify this provision and reemphasize the importance of a "buy not build" approach to procuring information technology systems by the federal government, including those that integrate AI.*²⁹

Custom-built technology solutions, whether developed by a federal agency or a government contractor, have historically had a high failure rate, are often more expensive than commercial alternatives, are costly and difficult to maintain, and quickly become obsolete.³⁰ Commercial solutions, including Workday's, are better positioned to help agencies streamline and standardize administrative processes, gain unprecedented visibility across organizations, and adapt quickly to change, including changes in fast-evolving technologies such as AI.

Existing federal acquisition rules and the *Federal Acquisition Streamlining Act of 1994*³¹ have long required the federal government to prioritize and procure commercial items, to the maximum extent practicable, over custom development.³² This should be no different when considering the procurement of secure, responsible AI tools. Given the fast-evolving nature of AI technology, as well as efforts by companies like Workday to develop AI in a responsible manner, it is critical that agencies turn to commercial solutions and expertise. We encourage OMB to make this clear in the final Memorandum.

OMB rightly identifies workforce challenges as a barrier to AI adoption.

As noted above, the federal government faces workforce challenges that are not limited to AI. At the same time, computer and AI literacy are increasingly needed for today's jobs, which we expect to grow in the future. We therefore applaud OMB for identifying workforce challenges as a barrier to agency adoption of AI. Consistent with the *Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence* (AI Executive Order), the Draft Memorandum highlights the need for the federal government to recruit, attract, and retain AI talent from across the country. Removing these barriers is paramount for the federal government to achieve *both* its AI and workforce goals. The emerging focus by the federal

²⁸ "Artificial Intelligence Index Report 2023." Stanford University. 2023.

https://aiindex.stanford.edu/wp-content/uploads/2023/04/HAI_AI-Index-Report_2023.pdf.

²⁹ Our recommendation is not intended to suggest that agencies should refrain from developing internal AI talent. This suggestion, instead, focuses on the "buy, not build" dichotomy and the availability of commercial items in the AI marketplace.

³⁰ Yaraghi, Niam. "Doomed: Challenges and solutions to government IT projects." August 25, 2015. Brookings. <https://www.brookings.edu/articles/doomed-challenges-and-solutions-to-government-it-projects/>.

³¹ P.L. 103-355.

³² U.S. Government Accountability Office. "Procurement Reform: Implementation of the Federal Acquisition Streamlining Act of 1994." Procurement Reform: Implementation of the Federal Acquisition Streamlining Act of 1994 | U.S. GAO, <https://www.gao.gov/products/t-nsiadaimd-95-190>.

government on skills-based hiring, rather than solely on traditional credentials, is the right lens through which to view these workforce priorities.³³

AI is unique in that it will drive change in the workplace and power the tools that workers and employers need to successfully navigate those changes. In many ways, a skills-based approach to talent is difficult to scale at a level necessary for the federal government without AI-driven tools implemented in a responsible manner. For example, the AI tools we deliver to our customers help people to make more informed decisions by surfacing new insights, identifying opportunities for career development, and improving workers' day-to-day by simplifying labor-intensive tasks.³⁴

Our guiding principle is that AI should be used in ways that augment, rather than displace people. We've found that this is an approach that builds trust with our customers and end users.³⁵ As such, Workday provides tools that enrich—but don't automate—human decisions and can make a skills-based approach to talent a reality. We encourage OMB to work with federal agencies to view AI as a tool to help address federal workforce challenges, including those being driven by federal AI-adoption.

Managing Risks from the Use of Artificial Intelligence

Impact assessments are the right AI accountability tool.

Workday strongly supports the use of impact assessments in AI governance frameworks and commends OMB for directing agencies to use them for "rights-impacting" AI systems. In 2021, we released a whitepaper, *Building Trust in Artificial Intelligence and Machine Learning*, in which we introduced a "Trustworthy by Design" framework.³⁶ Our whitepaper proposed that organizations adopt AI governance frameworks and conduct an AI impact assessment for each high-risk AI system that they intend to develop or deploy. We also called for developers to supply deployers with system information identifying the AI tool's intended purpose and any known limitations.

The state of AI technology and governance has advanced significantly since 2021. We nonetheless believe that impact assessments retain three important advantages over other potential AI accountability tools because:³⁷

- **Impact assessments are familiar and reliable.** Widely used in the fields of privacy and data protection, organizations are experienced at conducting impact assessments. The

³³ "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence." The White House, October 30, 2023.

<https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

³⁴ Chakraborty, Sayan. "Workday's Vision for AI." Workday Blog, March 10, 2023.

<https://blog.workday.com/en-us/2022/workdays-vision-for-ai.html>.

³⁵ Chakraborty, Sayan. "How AI and ML Are Powering the Future of Work." Workday Blog, July 6, 2023.

<https://blog.workday.com/en-us/2023/how-ai-and-ml-are-powering-future-work.html>.

³⁶ Building Trust in AI and ML through principles, practice, and policy, n.d.

<https://www.workday.com/content/dam/web/en-us/documents/whitepapers/building-trust-in-ai-ml-principles-practice-policy.pdf>.

³⁷ "Impact Assessments: Supporting AI Accountability & Trust." Access Partnership. 2023.

<https://www.workday.com/content/dam/web/en-us/documents/legal/access-partnership-workday-impact-assessment-paper.pdf>.

European Union’s *General Data Protection Regulation* and at least ten U.S. state privacy laws require private sector organizations to carry them out.³⁸ Under the *E-Government Act of 2002*, federal agencies are required to conduct privacy impact assessments before using a system that collects, maintains, or disseminates information about members of the public.³⁹ These serve as “one of the most valuable tools federal agencies use to ensure compliance with applicable privacy requirements and manage privacy risks.”⁴⁰

- **Impact assessments are practical** because they do not rely on technical standards or quantitative metrics, which in the field of AI governance are nascent.⁴¹ An impact assessment can function as a holistic risk assessment that accounts for an AI tool’s objective, design, intended use, and limitations.
- **Impact assessments are future-proof** because they can adapt as AI systems and AI governance evolve. With the launch of NIST’s new AI Safety Institute, there is growing focus on the need to develop “a new measurement science that will enable the identification of proven, scalable, and interoperable techniques and metrics” for trustworthy AI.⁴² As these metrics and international standards continue to develop, they can be incorporated into future impact assessments conducted by federal agencies.

For these reasons, there is a growing consensus among lawmakers, business leaders, and civil society that impact assessments for high-risk AI tools are the most promising accountability tool available today. *Accordingly, we agree with the proposed requirement, starting on August 1, 2024, that agencies complete AI impact assessments before using existing or new safety-impacting or rights-impacting AI.*

OMB should strengthen and streamline the proposed AI impact assessment process.

Workday makes the following recommendations to strengthen federal agencies’ use of AI impact assessments. First, impact assessments should focus on identifying, documenting, and mitigating “reasonably foreseeable” risks of harm to an individual’s rights or safety. The Draft Memorandum’s current language on “potential risks” invites federal agencies to speculate on *any possible* risks that can arise from the use of an AI tool, which may militate against a tool’s expected benefits altogether. Importantly, emerging benchmarks for AI governance in California and among key U.S. trading partners recognize “reasonably foreseeable” as a legal standard for mitigating harms.⁴³ The FTC has also issued guidance to businesses developing or using AI

³⁸ “Impact Assessments - a key part of AI accountability.” BSA, <https://www.bsa.org/policy-filings/impact-assessments-a-key-part-of-ai-accountability>.

³⁹ P.L. 107–347.

⁴⁰ FPC.Gov. Federal Privacy Council, <https://www.fpc.gov/elements-of-federal-privacy-program/privacy-impact-assessments/>.

⁴¹ The quality of an impact assessment, like all governance processes, is improved by technical standards. Unlike other accountability tools, however, technical standards are not *required* for impact assessments, which offer a holistic and robust identification, document, and mitigation of AI risk.

⁴² U.S. Artificial Intelligence Safety Institute. NIST. 2023. <https://www.nist.gov/artificial-intelligence/artificial-intelligence-safety-institute>.

⁴³ California’s *AB331* requires developers and deployers to conduct impact assessment to identify and mitigate the “reasonably foreseeable risk of algorithmic discrimination.” Canada’s draft *Artificial Intelligence and Data Act* and the European Union’s *Artificial Intelligence Act* are also both expected to integrate “reasonably foreseeable risk” as a legal standard for assessing and mitigating potential harms.

and other high-risk technology to “know the reasonably foreseeable risks and impact” of their products.⁴⁴

Second, OMB should require developers providing AI to the federal government to supply “data sheets” or “data cards” concerning training data. AI developers’ representations concerning the provenance, relevance, breadth, reliability, and quality of their datasets would better ensure that agencies have the information necessary to complete reporting requirements, without unduly delaying agency impact assessments. To streamline the impact assessment process, we recommend the use of common forms across federal agencies, which will rely on information regarding the quality and appropriateness of data in a standardized fashion.

Third, OMB should require each covered agency to develop a plan to dedicate sufficient resources for conducting impact assessments, testing, and independent evaluations within 120 to 180 days. We note that agencies will need to apply appropriate resources to fulfill these requirements if they are to avoid bottlenecks for systems already in use, in development, or in procurement prior to and after August 1, 2024.

The Draft Memorandum should center federal AI risk management practices on the NIST AI Risk Management Framework.

Workday is pleased that the Draft Memorandum encourages federal agencies to incorporate the NIST AI RMF into their governance processes. *We nonetheless urge OMB to go further than recognizing it as a “best practice,” as the Draft Memorandum currently does, and instead require agencies to use the AI RMF.* Doing so would advance a number of AI policy and governance goals, including: (1) ensuring federal agencies, contractors, and vendors have a common baseline by which to understand, assess, and communicate AI risk management practices and controls; and (2) signaling to the U.S. private sector and U.S. trading partners that the AI RMF is *the* benchmark for mapping, measuring, managing, and governing AI risk.

In addition to promoting a consistent approach to AI risk management and U.S. leadership in responsible AI governance, agency adoption of the NIST AI RMF will set stable expectations in the federal marketplace. In contrast to other “best practices,” the NIST AI RMF was developed at the direction of Congress through an open, multi-stakeholder process with participation from academia, civil society, and the business community. There is bipartisan, bicameral support in Congress for its adoption by federal agencies,⁴⁵ including *The Federal AI Risk Management Act*.⁴⁶ These proposals echo the National AI Advisory Committee’s (NAIAC) Year 1 Report to the President, which recommends federal use of the AI RMF.⁴⁷ A recent follow-up recommendation from the NAIAC outlines how the AI RMF can be implemented by federal agencies with a

⁴⁴ Atleson, Michael. “Keep your AI claims in check.” FTC. February 27, 2023. <https://www.ftc.gov/business-guidance/blog/2023/02/keep-your-ai-claims-check>.

⁴⁵ “Reps. Lieu, Lofgren and Stevens Urge Biden Admin to Require Federal Agencies to Adopt NIST AI Framework.” Office of Congressman Ted Lieu. <https://lieu.house.gov/media-center/press-releases/rep-lieu-lofgren-and-stevens-urge-biden-admin-require-federal-agencies>.

⁴⁶ “Sens. Morgan, Warner Introduce Legislation to Establish AI Guidelines for Federal Government.” Office of Senator Jerry Morgan. November 2, 2023. <https://www.moran.senate.gov/public/index.cfm/news-releases?id=17018C56-4B23-4DEC-A7B7-E899CD43C2C0>.

⁴⁷ “National Artificial Intelligence Advisory Committee Year 1 Report.” NAIAC, 2023. <https://www.ai.gov/wp-content/uploads/2023/05/NAIAC-Report-Year1.pdf>.

rights-respecting approach.⁴⁸ This underscores that the AI RMF outlines a robust, flexible approach to AI governance that can advance a key objective of the Draft Memorandum: the upholding of civil rights, civil liberties, and consumer protection. Accordingly, the AI RMF warrants further consideration as a required baseline for federal agencies than currently envisioned in the Draft Memorandum.

Federal agencies' independent evaluations should avoid premature third-party audit requirements.

The Draft Memorandum envisions federal agencies conducting an independent evaluation of impact assessments and testing results. Given resource constraints, federal agencies may ultimately contract these independent evaluations to third-parties. Although internal agency evaluations are appropriate, we caution OMB that evaluations that rely on third-party auditors may not be fit for purpose.

Although promising in the medium- and long-term, AI auditing is a field that is still in development, as there are neither consensus technical standards nor a common set of criteria to audit AI tools against. Notwithstanding meaningful steps in favor of professionalizing the AI governance field, there are currently no binding standards on third-party AI auditors specifically, which are necessary for ensuring auditing quality and integrity.⁴⁹ As noted in the White House's *2023 National AI Research and Development Strategic Plan*, absent such standards, the scalability of AI auditing is a significant practical challenge.⁵⁰ Without such building blocks, premature third-party audit requirements may unintentionally diminish trust in AI used in the federal enterprise by failing to promote consistent accountability and delaying the implementation of critical AI tools.

As one prominent researcher has noted, "inadequate audits or those without clear standards provide false assurance of compliance with norms and laws, 'audit-washing' problematic or illegal practices."⁵¹ Only once critical building blocks are in place can third-party AI audits serve as an important accountability tool to complement impact assessments.

The definition of "right-impacting" systems should be refined.

The Draft Memorandum appropriately adopts a risk-based approach to AI governance that maximizes the benefits of AI and minimizes the risks of potential harm. As key U.S. trading partners are demonstrating, including Canada, the European Union, the United Kingdom, and Singapore, a risk-based approach means applying rules to contexts and uses of AI that carry the highest risk of potential harm to individuals. To implement a risk-based approach,

⁴⁸ "Recommendations: Implementing the NIST AI RMF with a Rights Respecting Approach." The National Artificial Intelligence Advisory Committee. October 2023. https://ai.gov/wp-content/uploads/2023/11/Recommendations_Implementing-the-NIST-AI-RMF-with-a-Rights-Respecting-Approach.pdf.

⁴⁹ Hughes, Trevor. "The Time to Professionalize AI Governance is Now," IAPP. October 2, 2023. <https://iapp.org/news/a/the-time-to-professionalize-ai-governance-is-now/>.

⁵⁰ National Artificial Intelligence Research and Development Strategic Plan. The Executive Office of the President. 2023. <https://www.whitehouse.gov/wp-content/uploads/2023/05/National-Artificial-Intelligence-Research-and-Development-Strategic-Plan-2023-Update.pdf>.

⁵¹ Ellen P. Goodman and Julia Trehu, "AI Audit-Washing and Accountability," GMF. November 2022. <https://www.gmfus.org/sites/default/files/2022-11/Goodman%20%26%20Trehu%20-%20Algorithmic%20Auditing%20-%20paper.pdf>

policymakers have defined high-risk systems by focusing on: (1) the *purpose or intended purpose* for which an AI tool is used; and (2) *how* the AI tool's output is used to advance that purpose. Only by properly scoping *both* elements can a risk-based approach be workable in practice.

Although the Draft Memorandum's initial approach is promising, we recommend that OMB further refine its definition of a "rights-impacting" system. The Draft Memorandum currently sets a presumption for federal agencies that an AI tool is "rights-impacting" if it is (1) used in a certain set of decisions, activities, or contexts (2) to "control or meaningfully influence" outcomes. Workday urges OMB to refine this definition because, as written, it is likely to capture AI tools that are not high-risk, thereby undercutting the benefits of a risk-based approach to AI governance and slowing AI adoption in the federal enterprise.

First, we recommend that OMB focus solely on consequential *decisions*, rather than a broader set of *contexts* or *activities*. Consequential decisions in the employment context are determinations as to an individual's access to important life opportunities, such as the decision to hire, terminate, set pay, or promote an employee. As written, the Draft Memorandum captures a wide spectrum of potential AI use cases with no indication of how they are "rights-impacting." It is unclear, for example, how *all* "virtual or augmented reality workplace training programs" impact the rights of an individual. By contrast, an individual's *selection* into such a training program directly impacts their access to that opportunity. Unfortunately, the Draft Memorandum's use of undefined terms (e.g., "workplace management system") adds to this uncertainty and invites different and likely competing interpretations by federal agencies. Workday urges OMB to set an exhaustive and clearly defined list of consequential decisions. Doing so will provide much needed clarity about which AI use cases are presumed to be "rights-impacting."

Second, we recommend that OMB focus the Draft Memorandum's scope on AI tools that *control* outcomes. This will improve the definition of "rights-impacting" systems in two respects. It will address the ambiguities associated with the current language on "control or meaningfully influence," because "meaningfully influence" is an unclear threshold.⁵² Importantly, it will also distinguish between fully automated AI tools and those with an informed human-in-the-loop in the decision.⁵³ An AI tool that *replaces* human decisions poses a higher risk than an AI tool that *enriches* them. This is because an automated tool can make consequential decisions at a higher volume and velocity than a human, and without the judgment that a human brings. By contrast, when an informed human is in the loop, they can leverage an AI tool's insights and remain in control of and accountable for the final decision. The Draft Memorandum already recognizes that a human-in-the-loop approach presents a lower-risk than fully automated ones: its definition for "rights-impacting" should recognize this as well. Agencies that use "rights-impacting" systems, for example, would be required to have ongoing human training and assessment and human consideration in final decisions. By focusing on AI tools that are fully automated and scoping out those with an informed human-in-the-loop, a revised definition of "rights-impacting" AI would focus agencies' limited resources on tools with the highest risk profiles.

⁵² If OMB believes a definition of "control" is needed, we recommend "the factor that is weighted more than all other factors together."

⁵³ Earlier this year, Workday surveyed 1,000 senior decision-makers in human resources, finance, and technology for their perspectives on AI. Nearly all (93 percent) said it was important to keep a "human in the loop" when making significant decisions. "AI IQ: Insights on AI in the Enterprise." Workday. June 2023. <https://www.workday.com/en-us/why-workday/our-technology/artificial-intelligence/research/ai-iq.html>.

OMB should encourage consistency in “rights-” and “safety-impacting” AI determinations by distinguishing between agency-specific AI systems and enterprise AI systems.

The Draft Memorandum treats the designation of “right-impacting” and “safety-impacting” AI as *presumptions* and enables federal agencies to either waive or make additional designations. This approach will likely result in inconsistent designations of which AI tools are or are not high-risk throughout the federal government. While agencies may be best placed to designate AI tools as “rights-” or “safety-impacting” if the use case relates to their mission, it is unclear how different designations of *enterprise* AI solutions advances AI risk management.

Workday recommends that OMB distinguish between agencies’ mission-specific AI systems (for which individual agencies are best equipped to tailor definitions and make determinations) and enterprise systems such as human resources and financial management systems (for which a uniform determination across the whole of government is more desirable). This will promote consistency in determinations of “rights-” and “safety-impacting” AI tools throughout the federal government and encourage a more consistent level of protection.

In the case of enterprise systems that may serve multiple agencies, CAIOs should, to the maximum extent practicable, coordinate through the interagency council that will be established under Section 10.1(a) of the *AI Executive Order*, to ensure consistent determinations about “rights-” or “safety-impacting” AI systems and features. Coordination through the interagency council would benefit AI adoption across the federal government by reducing duplicative agency efforts and compliance burdens on federal contractors and vendors.

We also recommend that OMB direct CAIOs to distinguish between mission-specific and enterprise AI uses in agencies’ annual AI use case inventories. This information would enable OMB to catalog widely-used AI features for which government-wide determinations would be most valuable, and avoid inconsistent determinations between multiple agencies utilizing the same enterprise AI features for the same or similar purposes. As discussed in further detail below, the interagency council should also develop and implement frameworks for mutual recognition of AI authorizations and monitoring activities, and to ensure that enterprise AI systems are not subject to unnecessarily duplicative, wasteful, or disparate requirements.

OMB Should Enhance Standardization and Encourage Reciprocity by Minimizing Agency-Specific Definitions and Practices.

As noted above, adoption of beneficial AI systems by the federal government is best served by uniform definitions and practices across federal agencies. An agency-by-agency approach to developing key definitions and requirements will likely result in fragmentation that will subject a single AI tool to diverging authorization, testing, and monitoring practices between agencies. This would result in an expensive repetition of effort for both government agencies and AI developers without a corresponding benefit.

To avoid this result, OMB should remove language in the Draft Memorandum that encourages agencies to develop agency-specific definitions and practices. We recommend OMB instead establish a framework for standardized and reciprocal AI authorization and monitoring processes. The current FedRAMP program provides a model on which this approach could be designed, emphasizing reciprocity of AI authorizations across agencies and a presumption that the authorization and monitoring activities of an initial sponsoring agency are adequate for

similar use cases at any other agency.⁵⁴ In short, a single set of definitions and required practices would enable an efficient AI authorization framework inspired by the FedRAMP program and would enable government-wide AI adoption. Agency-specific deviations would inhibit this effort.

To the extent that agencies need the flexibility to identify agency-specific purposes, definitions, or practices, we recommend that they be limited to only those AI systems designed to address a single agency's mission. Workday also recommends that agency-specific deviations from OMB guidance be applied to AI systems only after the agency policies are clearly defined and made subject to public comments via Notice in the Federal Register.

AI Governance Boards will strengthen federal agencies' AI risk management.

We applaud the Draft Memorandum's requirement that federal agencies institute an AI Governance Board. As discussed above, Workday has instituted a Responsible AI Advisory Board to ensure senior leadership and provide executive oversight of our responsible AI program.⁵⁵ This internal advisory group meets regularly to review and approve new aspects of Workday's responsible AI program, and to advise on novel issues that are not currently contemplated by existing policies or processes—helping to address edge cases and escalations as needed. The advisory board members' seniority and cross-disciplinary diversity enables it to make difficult decisions, mandate necessary risk mitigations, and delay or stop product releases if the technology fails to align with the company's AI principles and core values.

OMB should clarify the Draft Memorandum's proposed opt-out requirement.

The Draft Memorandum would require agencies deploying “right-impacting” AI tools to “provide and maintain a mechanism to conveniently opt out from AI functionality in favor of a human alternative.” Given the broad definition of “rights-impacting” AI, we encourage OMB to clarify that opt out rights apply to AI tools used for consequential *decisions*, where the tool's output is the *sole* criterion for the decision. This would bring the Draft Memorandum in line with commonly accepted legal requirements, including the EU's *General Data Protection Regulation*, which recognize that an opt out requirement is appropriate for fully automated systems, rather than those with a human-in-the-loop. Moreover, as AI features are being integrated throughout commercial software, a broad opt out from all AI tools will likely serve as a barrier to AI adoption in the federal enterprise.

Managing Risks in Federal Procurement

OMB Should Clarify the Definition of “AI Contract” and Exclude “Common Commercial Products.”

To ensure consistency with section 7224(d) of the *Advancing American AI Act* and Section 10.1(d)(ii) of the *AI Executive Order*, Workday recommends that the Draft Memorandum adopt the definition of “artificial intelligence system” set forth in section 7223(4) of the Act:

⁵⁴ See e.g. 44 U.S.C. § 3613 (“The assessment of security controls and materials within the authorization package for a FedRAMP authorization shall be presumed adequate for use in an agency authorization to operate cloud computing products and services”).

⁵⁵ “Empowering Innovation through Responsible AI Governance.” Workday Blog. <https://forms.workday.com/en-us/whitepapers/empowering-innovation-through-responsible-ai-governance/form.html>.

ARTIFICIAL INTELLIGENCE SYSTEM.—The 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.

We suggest that OMB issue additional definitions and guidance to enable covered agencies to distinguish “AI contracts” within the meaning of the Draft Memorandum (or “contracts for the acquisition of an artificial intelligence system or service” within the meaning of the section 7224(d) of the Act), from contracts for any “common commercial product” excluded from the definition of AI system by section 7223(4)(B) of the Act. Given procurement preferences in favor of commercial products and services “consistent with customary commercial practice,” this guidance should refrain from applying new government-specific contract requirements to the acquisition of common commercial products—especially commercial items and commercial-off-the-shelf software products—that embed AI.

OMB should maximize the value of data for AI consistent with commercial terms and existing data rights clauses.

When a covered agency intends to procure an “artificial intelligence system or service” within the meaning of the section 7224(d) of the Act, the procurement laws and regulations applicable to commercial products and services still apply. With respect to data rights, we believe that covered agencies should not seek additional rights to data that are inconsistent with commercial practice or beyond the requirements of existing FAR data rights clauses.⁵⁶

For example, customary commercial practices and FAR data rights clauses would not grant the government buyer rights in data first produced outside of the performance of the contract, such as rights in a dataset initially used to train an AI model. Nor would they grant the government rights in model improvements that are a part of commercial computer software provided under the terms of a commercial license. Standard FAR data rights clauses and customary commercial practices would, however, grant the government rights in technical data first produced in the performance of a contract, such as rights in improvements or enrichments of government data developed through data cleaning or data labeling efforts. No further regulatory requirements are necessary to protect these government rights.

⁵⁶ FAR § 52.227-14; Defense Federal Acquisition Regulation Supplement § 252.227-7015.

Where rights in data produced under a contract will facilitate the government's continued design, development, testing, and operation of AI, we recommend that OMB not encourage agencies to seek additional rights that are inconsistent with commercial terms. Similarly, while agencies may consider provisions that would protect government data from use to train or improve the functionality of commercial AI without express permission from the agency, agencies should also understand that limiting the use of data – especially de-identified, disaggregated, anonymized data – in an AI developer's general product improvement cycle may be inconsistent with commercial practices.

OMB Should prepare acquisition planning materials and contract clauses for procurement of “rights-” and “safety-impacting” AI and encourage bilateral modification of impacted contracts.

While the minimum practices for “rights-” and “safety-impacting” uses of AI are intended to be agency responsibilities, the practices applied to an AI developer's systems will inevitably impact the scope of the contractor's efforts under an AI contract. It is therefore critical that agencies identify potential for “rights-” or “safety-impacting” uses of AI as early as possible in acquisition planning, and include requirements in their solicitations and contracts. OMB should assist agencies in the effort by preparing acquisition planning checklists and other guidance for identifying safety- or rights-impacting uses of AI – especially in “non-AI” procurements. OMB should also prepare sample contract requirements or clauses for agencies to include in their solicitations.

Finally, OMB should recognize the importance of minimizing disruptive changes to existing and future contracts that did not anticipate the requirements of the minimum practices at the time that they were executed by the parties. We suggest that the specific requirements for an AI system developer to support the minimum practices be included in the solicitation, and the contract for the AI product or services.

In the event that the requirements were not included because the contract was executed prior to the Draft Memorandum, or because the product or service was not identified as “rights-” or “safety-impacting” AI at the time, additional obligations should not be imposed on the contractor except by bilateral modification of the contract. This guidance would support contractors' certainty in their obligations as the government continues to develop frameworks for the use of AI, and insulate contractors from shifting determinations or waivers that may result from changes in an agency's CAIO, or updates to policies and practices over time. This will also reduce any subsequent litigation between the contractors and the agency relating to any such modifications.

Conclusion

Workday appreciates OMB's efforts to produce a Draft Memorandum that recognizes the benefits of federal agencies adopting responsible AI. The challenges facing the federal government and the importance of agencies' missions are only growing, making the availability of modern and secure cloud-based systems and AI tools essential.

At the same time, AI adoption can result in unintended consequences. Workday strongly believes in the need for smart AI guardrails and supports regulations that ensure AI can be deployed safely, users can trust the technology, and developers can continue to innovate. We

congratulate OMB on taking a thoughtful approach and embracing a risk-based path forward that uses impact assessments.

We offer the above comments, which are informed by our approach to responsibly developing AI tools and our engagement with policymakers on AI issues at the global, federal, state, and local levels, to strengthen OMB's proposal. Please contact Evangelos Razis, Senior Manager, Public Policy, (evangelos.razis@workday.com) if there is additional information we can provide.