Workday's Response to the National Institute of Standards & Technology's Request for Information on an *Artificial Intelligence Risk Management Framework*

August 2021

Workday is pleased to respond to the National Institute of Standards and Technology's (NIST) request for information on an *Artificial Intelligence* (AI) *Risk Management Framework* (RMF).

Workday is a leading provider of enterprise cloud applications for finance and human resources, helping customers adapt and thrive in a changing world. Workday applications for financial management, human resources, planning, spend management, and analytics have been adopted by thousands of organizations in the U.S. and around the world and across industries—from medium-sized businesses to more than 45 percent of the Fortune 500. Workday incorporates machine learning (ML) technology within our applications that enable customers to make more informed decisions and accelerate operations, as well as assist workers with data-driven predictions that lead to better outcomes. Workday believes ML technology has the potential to impact enterprises in the near-term by making operations more efficient. In the longer term, organizations will be able to reorganize their operations around machine learning and AI's unique possibilities.

AI is becoming an ever-increasing and transformative presence in our lives, driving human progress in countless ways. To achieve AI's full potential, however, there must be broad confidence that it is being developed ethically and used responsibly. With this in mind, Workday welcomes NIST's efforts to develop "forward-thinking approaches that support innovation and confidence in AI systems" and is pleased to offer the following comments.

I. Support for NIST's AI Risk Management Framework

Workday is a firm supporter of NIST's efforts to advance trustworthiness in AI, including through the development of a risk management framework. The issue of AI trustworthiness is ready-made for NIST, as it has a well-developed track record in convening government, industry, and other stakeholders to cooperatively develop cutting-edge voluntary frameworks. Given the success of the *Framework for Improving Critical Infrastructure Cybersecurity* and the *Privacy Framework – An Enterprise Risk Management Tool*, NIST has best-in-class experience with open, transparent, and collaborative processes.

Workday has played an active role in supporting NIST's leadership in this area through extensive legislative efforts and coalition building. We were pleased to see AI framework provisions receive bipartisan support and inclusion in the *Fiscal Year 2021*

National Defense Authorization Act. We are keen to build on these efforts by constructively engaging in NIST's AI workstreams.

As NIST, together with stakeholders, embarks on the path of developing an AI RMF, timeliness is essential. The European Union is moving forward with the *Artificial Intelligence Act* (AI Act), which would require organizations developing and deploying covered AI applications to implement risk-management systems. Workday recently provided comments to the European Commission on this consequential legislation. It is essential that the U.S., EU, and other like-minded governments harmonize their approaches to AI risk management and regulation so as to prevent unneeded barriers to transatlantic trade, investment, and innovation from developing. Indeed, NIST's RMF may facilitate enterprises' compliance with their future legal obligations under the AI Act. With Europe's proposal certain to become law, such harmonization is impossible if the U.S. lacks a consensus approach to managing AI risks, voluntary or otherwise. NIST's work on an AI RMF is very likely the most expeditious vehicle for establishing a consensus U.S. approach to building trustworthy AI.

If done in a timely, collaborative, and iterative manner, NIST's work will serve as an important U.S. contribution to global AI policy and to international regulatory cooperation, including to discussions under the auspices of the U.S.-EU Trade and Technology Council. These efforts would build on U.S. leadership at the Organization for Economic Cooperation and Development (OECD) and the Global Partnership on Artificial Intelligence, which bolster U.S. economic competitiveness and values.

II. Workday's Trustworthy AI Framework

In addition to NIST's work on an AI RMF, Workday supports AI regulation that is risk-based, enables innovation, and meaningfully addresses issues of trustworthiness. Earlier this year, we released a white paper, *Building Trust in Artificial Intelligence and Machine Learning*, outlining a framework for promoting trust, accountability, and transparency, while also giving organizations broad flexibility to innovate. This paper offers an informed perspective based on our experience in offering ML-based services to our business customers as NIST begins the framework development process.

Specific provisions we call for in our "Trustworthy by Design" regulatory framework would require organizations to:

- *Adopt Principles & Publish a Trustworthy AI Policy.*
  The paper calls for organizations to adopt principles setting out their trustworthy AI commitments and publish a public trustworthy AI policy addressing identified core elements of ethical artificial intelligence. Organizations would be expected to provide only a summary of their governance framework in the trustworthy AI policy itself. It would serve as a reference point for regulators and allow an appropriate agency to impose sanctions for misrepresentations.

- *Adopt an AI Governance Framework.*
  The governance framework should include:
  - Designation of a senior leader and establishing a trustworthy AI compliance team;
  - An approach to AI Impact Assessments and data documentation; and
  - Commitments for personnel training and providing cross-company compliance resources.

- *Implement Procedures to Identify and Mitigate Harmful Bias.*
  The paper calls on AI developers to implement and describe procedures to identify and mitigate sources of potentially harmful bias in their AI systems, both in the AI models they develop and in the data they use (including data used to train their systems and the data analyzed by those systems in real-world settings). It also calls on AI developers to document the procedures they use to test for, identify, and mitigate the effects of potentially harmful bias, as well as establish diverse teams to design and develop AI systems.

- *Conduct Impact Assessments.*
  The paper recommends a separate impact assessment for each AI system that developers create. That impact assessment should identify potential risks to individuals or society, reasonably quantifying the amount of risk and degree of potential harm and listing safeguards adopted to mitigate these risks to an acceptable level.

- *Maintain Data Documentation.*
  The paper calls for AI developers to document the provenance of AI training data and take reasonable steps to test whether the use of these datasets may lead to unfair or discriminatory outcomes.

- *Provide User Transparency & Recourse.*
  The paper calls for AI deployers to provide users appropriate transparency about the individual AI systems they interact with, the safeguards implemented against untrustworthy uses of the system, and the recourse available to them, such as the ability to appeal decisions to a person. In the majority of instances, it is the AI deployer who has the most direct relationship with affected individuals and is therefore the actor best suited communicate with them.

- *Supply Information for Deployers.*
  The paper calls for AI developers to provide AI deployers with:
  - The intended purpose and the acceptable use of the AI system;
  - Steps on how the system can be properly deployed; and
  - Any known limitations in the system, model notices, and any unintended or unacceptable uses, as well as the level of human oversight, if any, that deployers should provide.

- *Provide Deployer Support for Individuals.*
  The paper calls for AI developers to provide AI deployers with sample notices and explanations, which deployers would use to communicate key aspects of the system to affected individuals. These explanations also might include:
    - High-level description of the internal workings of the AI system; and
    - The logic of the AI system and/or information about the accuracy, reliability, safety, or other features of the system.

Workday's regulatory framework would require companies to publish their trustworthy AI policies and enable customers and users to compare policies between enterprises. This in turn would create market-based incentives for companies to adopt robust, meaningful policies. The framework gives AI producers flexibility to adopt principles and practices that are most appropriate for their businesses, tailored to the type and degree of risk their AI systems present. As previously mentioned, it also envisions that future standards and industry best practices would be developed that will play a key role in enabling organizations to demonstrate trustworthiness.

There is an emerging and broad consensus throughout industry and stakeholders in the U.S. and abroad around the baseline principles that should guide trustworthy AI. These baseline principles, which NIST should account for, include fairness, transparency, accountability, and respect for fundamental human rights. For example, similar to Workday's AI regulatory framework, BSA | The Software Alliance recently published their *Confronting Bias: BSA's Framework to Build Trust in AI,* a detailed AI bias risk management framework that organizations can use to perform impact assessments to identify and mitigate risks of bias that may emerge throughout an AI system's lifecycle.

III. Response to Specific Requests for Feedback

A. Challenges to Managing AI Risk

AI risk management faces a number of challenges that reflect both the breadth of AI use cases and the emerging nature of the field. The AI ecosystem is home to a multitude of diverse stakeholders, including developers, users, deployers, and, of course, consumers. For enterprises, identifying, assessing, prioritizing, responding to, and communicating risks across complicated business relationships at scale is no small task. These challenges are amplified by the heterogeneity of AI risks, whether in degree or in kind, which include harmful bias, health and safety, privacy, and consumer protection, among others.

While the field of AI risk management is still maturing, NIST's cybersecurity and privacy frameworks serve as useful examples for how to address these challenges. An AI RMF can serve as a common grammar that organizations can use to communicate their risk management practices. With both Workday and BSA emphasizing the need for flexibility, we urge NIST to recognize that any proposed path forward that seeks to

promote AI ethics and trust across millions of scenarios and use cases in a prescriptive, one-size-fits-all manner will be unworkable.

B. Relevant Frameworks, Principles, & Policies

It is worth recognizing that, in contrast to prior NIST frameworks, the area of AI is comparatively less mature in terms of policy, regulation, standards, and best practices. Where past NIST frameworks on privacy and cybersecurity drew on more mature bodies of work, the AI RMF will be developed in an actively emerging field, without well-established approaches for systematically governing AI risks and its uses. In particular, no existing U.S. regulation holistically addresses the AI issues posed by the AI RMF, and many technical standards and best practices are still in early development. Yet this reality, in fact, makes the AI RMF exercise even more critical, because it is likely to serve as the groundwork for future approaches for governing AI in the U.S., including eventual regulation.

While AI regulation remains nascent in the U.S., the AI RMF should account for AI policy initiatives emerging elsewhere in the world. The EU's proposed AI Act, for example, moves beyond voluntary approaches to directly regulate AI applications deemed as high-risk. As Europe is the U.S.'s largest bilateral trade and investment partner, we recommend NIST consider how to leverage the AI RMF to support transatlantic regulatory cooperation, an imperative supported by political leaders at June's U.S.-EU Summit. Additionally, the Government of Singapore's *Model AI Governance Framework* is relevant, both for its consideration of governance programs and as a risk-based framework developed in collaboration with stakeholders through an iterative process.

Against this backdrop, NIST should ensure the AI RMF is iterative, scalable, and able to effectively incorporate and build on new regulations, best practices, and technical standards as they are developed.

C. Inclusion of Governance Issues

Workday strongly endorses the inclusion of governance issues in NIST's forthcoming AI risk-management framework. Simply put, organizations put in place governance programs to make tangible the goals, principles, and values underpinning trustworthiness. Absent such programs, AI risk-management is a constellation of tools and practices implemented unevenly, without transparency and accountability. Recognizing that the specifics of a governance program will necessarily vary according to the size and capacity of organizations, NIST should consider their basic elements. These include the involvement of senior management, such as appropriate C-Suite executives, to oversee the company's AI product development lifecycle, and a trustworthy AI compliance team responsible for carrying out impact assessments, documentation, training, and serving as a cross-company resource. In doing so, NIST's

AI RMF can assist enterprises and regulators alike by providing a standardized approach for ensuring accountability.

IV. Conclusion

Thank you for the opportunity to respond to NIST's request for information on an AI risk management framework. We congratulate NIST on the work put into AI thus far, including stakeholder involvement. Workday welcomes opportunities to support NIST in its efforts to develop a workable AI framework that is timely, impactful, and promotes trusted AI innovation and U.S. leadership on global AI regulatory cooperation.

We stand ready to provide further information and to answer any additional questions. Please do not hesitate to reach out to Evangelos Razis at evangelos.razis@workday.com for assistance.