Workday's Comments on NIST's Second Draft of the AI Risk Management Framework

September 2022

Workday welcomes the opportunity to comment on the National Institute of Standards & Technology's (NIST) second draft of the *Artificial Intelligence Risk Management Framework* (AI RMF or Framework).

Workday is a leading provider of enterprise cloud software applications for finance and human resources, helping our customers adapt and thrive in a changing world. Our applications have been adopted by thousands of organizations in the United States and globally—from medium-sized businesses to more than 50 percent of the *Fortune* 500. Workday incorporates machine learning into our software to enable customers to make more informed decisions and accelerate their operations, as well as to assist workers with data-driven predictions that lead to better outcomes.

Workday is a staunch supporter of NIST's pioneering efforts to develop the Framework. In our view, a strong foundation of trust is needed for artificial intelligence-based (AI) technologies to reach their full potential. These comments build on Workday's previous contributions to NIST's AI workstreams, including our participation in the first RMF workshop and our written comments on NIST's first draft, concept paper, and request for information.

I. General Comments

Workday is encouraged by NIST's progress in developing the Framework, including the release of the companion Playbook. As NIST seeks to publish a final first version of the AI RMF, we encourage it to consider what steps are needed to facilitate the timely and widespread adoption of the Framework by developers and deployers. Our comments below are intended to support NIST in achieving this goal.

II. Specific Comments

A. Impact Assessments

Workday welcomes the Framework's reference to impact assessments. Impact assessments are a tried-and-true way for organizations to document how they identify, test for, and mitigate the risks posed by technologies.

NIST correctly characterizes AI software testing standards as "underdeveloped," a fact which will hold back other accountability tools, such as AI auditing, in the near-term. As technical standards continue to develop, impact assessments represent a pragmatic way to advance trustworthiness at the level of the individual AI system. Impact assessments are widely used by companies for privacy and data protection purposes and are increasingly being adopted in AI governance programs.

We therefore recommend that NIST expand on how the AI RMF can be operationalized through impact assessments, which it may wish to do in the Playbook. In doing so, NIST would facilitate the adoption of the Framework by organizations, as many already use these accountability tools in their privacy programs.

B. Human-in-the-Loop

Workday notes the Framework's discussion of "human-in-the-loop" (HITL) with interest. As NIST expressed in its recent publication, *Towards a Standard for Identifying and Managing Bias in Artificial Intelligence*, no single practice is a panacea for managing AI risks. Nonetheless, HITL approaches can provide human oversight *and human accountability* for harms caused by AI systems in use, including when AI systems make consequential decisions.

The Framework highlights "unclear expectations" and "imprecise governance" as two "considerations" in HITL approaches. We encourage NIST to expand on how these considerations may be addressed so that the benefits of HITL may be realized. The Framework's transparency and accountability principles are relevant in this regard. Users of consequential AI systems, for example, should be equipped with sufficient information by their organizations and by developers so they can use systems as intended. Clear lines of responsibility and administrative and technical safeguards can also be put in place by deployers to ensure accountability.

C. Recognize Risks Arising from Deployment and Use

The Framework rightly characterizes AI risk management as a responsibility shared between developers and deployers. Yet the Framework focuses its attention on risks posed by AI systems from third-party vendors. While this emphasis is an important one, we recommend that NIST also expand on the risks associated with *deploying* those systems with insufficient administrative and technical safeguards.

An organization that develops an AI system is not always the same as the one deploying the system, and indeed may have limited visibility into how the system is being used. The way in which an AI system is implemented and used by an organization can create new risks, including risks that may not be foreseen by the system's developer. A developer, for example, may provide information to avoid the "Oracle problem" highlighted by NIST in the Framework. Yet a deployer may use the system in a manner that is not mindful of this information. Accordingly, we encourage NIST to offer a holistic account of AI risks consistent with its view that AI risk management is a shared responsibility.

III. Conclusion

Thank you for the opportunity to comment on NIST's second draft of the Framework. Please do not hesitate to reach out to Evangelos Razis at
if we can provide further information or answer any questions.