



September 28, 2022

US Department of Commerce
National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, MD 20899

Submitted via email: AIFramework@nist.gov

BSA | The Software Alliance Comments on NIST AI Risk Management Framework: Second Draft

BSA | The Software Alliance appreciates the opportunity to provide comments on the second draft of the AI Risk Management Framework (AI RMF) under development by the National Institute of Standards and Technology (NIST). BSA is the leading advocate for the global software industry before governments and in the international marketplace. Our members are at the forefront of software-enabled innovation that is fueling global economic growth and helping businesses of all sizes leverage the benefits of cloud computing and AI-enabled products and services.¹

Enterprise software, including artificial intelligence (AI), is accelerating digital transformation in every sector of the economy. AI is not just about robots, self-driving vehicles, or social media. It is used by businesses of all sizes to create the products and services they provide to consumers, to improve their internal operations, and to enhance their capacity to make data-informed decisions. BSA members are industry leaders that provide businesses in every sector of the economy with the trusted tools they need to leverage the benefits of AI. In particular, BSA members help businesses innovate and grow by providing cloud storage services, customer relationship management software, human resource management programs, identity management services, and collaboration software.

As leaders in the development of enterprise AI, BSA members have unique insights into the technology's tremendous potential to spur digital transformation and the policies that can best support the responsible use of AI. BSA's views are informed by our recent experience working with BSA member companies to develop the BSA Framework to Build Trust in AI,² a

¹ BSA's members include: Adobe, Alteryx, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, CrowdStrike, Databricks, DocuSign, Dropbox, Graphisoft, IBM, Informatica, Intel, Kyndryl, MathWorks, Microsoft, Okta, Oracle, Prokon, PTC, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Splunk, Trend Micro, Trimble Solutions Corporation, TriNet, Twilio, Unity Technologies, Inc., Workday, Zendesk, and Zoom Video Communications, Inc.

² See BSA, *Confronting Bias: BSA's Framework to Build Trust in AI*, available at <https://www.bsa.org/reports/confronting-bias-bsas-framework-to-build-trust-in-ai>.

risk management framework for mitigating the potential for unintended bias throughout an AI system's lifecycle. Built on a vast body of research and informed by the experience of leading AI developers, the BSA framework outlines a lifecycle-based approach for performing impact assessments to identify risks of AI bias and highlights corresponding risk mitigation best practices.

BSA has leveraged this experience to provide input on NIST's AI RMF. The AI RMF has the potential to establish a shared framework for identifying and mitigating risks throughout the AI system lifecycle and to enhance communication among multiple stakeholders involved in the development and deployment of an AI system. We note that the changes in the second draft of the AI RMF have brought the framework more in line with its intended goals.

Our comments below reiterate several points made in our comments on the initial draft, which are attached to this submission. We encourage NIST to address five key issues:

- expand the guidance on impact assessments;
- address the complexity of the AI ecosystem and the varying roles that stakeholders play in the development and deployment of AI systems;
- provide guidance on scenarios when organizations customize third-party software for deployment;
- clarify that use of the term "external stakeholders" includes other in-house personnel, and that the need to consult stakeholders outside of the organization will vary based on the risks involved; and
- clarify that the activities included in the descriptions of AI design and AI development do not include contextual decisions about the use of AI systems.

I. Impact Assessments

BSA applauds NIST's efforts to integrate the use of impact assessments into the framework. In addition to explicit mention of the usefulness of impact assessments in the document, we believe changes made to the Manage function also reflect valuable steps that can be taken as part of an impact assessment process. We encourage NIST to further develop the guidance on impact assessments in the AI RMF. For example, the framework could leverage existing NIST language on impact assessments outlined in the NIST publication *Towards a Standard for Identifying and Managing Bias in Artificial Intelligence*.³ In addition, by clarifying that the risk analysis under the AI RMF should focus on the impacts of an AI system and providing more clarity about how the trustworthiness characteristics of AI systems should be used as part of that holistic examination, the framework will better serve the needs of its target audience.

II. Roles of AI Developers and Deployers

We appreciate NIST's recognition that there is shared responsibility among the various actors throughout the AI ecosystem, but we encourage NIST to expand on this important point. For example, NIST could provide guidance to address circumstances where multiple stakeholders are involved in the development and deployment of AI systems and acknowledge that these responsibilities may be split across business units and between

³ NIST, Special Publication 1270, *Towards a Standard for Identifying and Managing Bias in Artificial Intelligence*, <https://www.nist.gov/publications/towards-standard-identifying-and-managing-bias-artificial-intelligence>.

organizations, particularly in instances where a company is providing a customizable AI system to a customer organization.

As currently drafted, the Framework Core appears to presuppose that the organization using the AI RMF will have full visibility into and control over the entire lifecycle of the AI system it is evaluating. For instance, while the Map function largely focuses on an analysis of a system's deployment context,⁴ it also calls for an examination of system artefacts that may only be available to the entity that trained the underlying model.⁵ The Framework Core provides little guidance about how organizations should utilize the AI RMF in circumstances where they are preparing to deploy an AI system or capability that may have been acquired from an external vendor.⁶ Similarly, the Framework Core provides little guidance to vendors who develop general-purpose AI systems, which may have limited insight into how their customers deploy the systems.

NIST should consider adding high-level guidance to Part 1 of the AI RMF to explain further that risk management will in many instances be a shared responsibility that encompasses multiple entities with differing roles and responsibilities within the AI ecosystem, including organizations that develop AI systems for use by other companies (i.e., AI Developers) and entities that deploy AI systems that they have acquired from external vendors (i.e., AI Deployers). While NIST should avoid drawing bright lines about how specific responsibilities should be assigned, it would be helpful for the AI RMF to acknowledge that the appropriate allocation of risk management will depend on the nature of the underlying model and the extent to which it may be customized and/or re-trained by the AI Deployer.⁷ Such guidance would serve as a useful tool for facilitating conversations between vendors of AI services and their customers to ensure that there is a shared understanding about their respective roles and responsibilities. Our comments on the initial draft further elaborate on this important issue.

III. Risks Related to Third Parties

The AI RMF focuses on third-party entities that are providers, developers, or vendors of data, algorithms, models and/or systems. However, its reference to third parties is focused primarily on the risk of use of third parties in AI systems.

We suggest adding language on the risks of companies deploying systems developed by a third party for uses that are either unintended or outside of the intended scope of uses for that AI system. The developer and deployer of an AI system may be different actors, and the

⁴ For example, the subcategories in the Map function call for an analysis of the "intended purpose," "settings in which the AI system will be deployed," and the "business value or context of business use."

⁵ For example, Map ID 2 includes a subcategory that is focused on considerations related to "experimental design, data collection and selection (e.g., availability, representativeness, suitability), and construct validation."

⁶ The Govern function does acknowledge the importance of maintaining policies to "address AI risks arising from third-party software and data and supply chain issues." But, beyond the suggested outcome to have contingency plans in the event of failure of third-party data or AI systems deemed to be high risk, the Framework Core lacks meaningful guidance about how to navigate these risks.

⁷ The OECD Framework for the Classification of AI Systems adopts a similar approach for assigning risk management responsibilities. See OECD Framework for the Classification of AI Systems 48, February 2022, <https://www.oecd-ilibrary.org/docserver/cb6d9eca-en.pdf?expires=1649808351&id=id&accname=guest&checksum=74B738F154B4F05D18B7B3D8B3477CE0>.

section on third-party risks would be improved by emphasizing that risk can emerge both from third-party software itself and from users of a third-party software customizing a product.

IV. References to External Stakeholders

The AI RMF frequently mentions external stakeholders as an important part of AI risk management. As we mentioned in our initial comments, we think it would be valuable to ensure that the use of the term “external stakeholders” throughout the document refers not only to third-party organizations, but also to in-house personnel. Specifically, it may be helpful to clarify that “external stakeholders” related to AI systems could also be developers or deployers outside of the team that originally developed the system. We also note that the need to consult other stakeholders may vary based on the risk of the AI system. We therefore suggest the following changes to the document:

- We suggest NIST adjust item 5.1 under “Govern” (page 20) to read: “Organizational policies and practices are in place to collect, consider, prioritize and integrate feedback from *stakeholders external to the team that developed or deployed the AI system* regarding the potential individual and societal impacts related to AI risks.”
- We suggest NIST adjust language in the second paragraph on page 21 regarding engagement with external stakeholders to read: “Implementing this function necessitates a broad set of perspectives from a diverse internal team and engagement with *stakeholders external to the team that developed or deployed the AI system. Engagement with external stakeholders may vary in necessity based on the risk level of a particular AI system and the makeup of an internal team, particularly when it comes to considerations of diversity and expertise.*”
- We suggest NIST adjust item 1.3 under “Measure” to acknowledge that the necessity of consultation with external stakeholders and affected communities as part of assessments may vary given differing levels of risk associated with an AI system.

V. Definitions of AI Design, AI Development, and AI Deployment

The AI RMF provides descriptions of AI actor tasks, including for AI design, development, and deployment. To ensure the roles and responsibilities are accurately reflected, it may be helpful to clarify that deployers are the entities that make decisions about how AI systems are used. Accordingly, we suggest that NIST add the following sentence to the AI Design description: “*AI actors in the design phase create the concept and objectives of AI systems, but deployers are responsible for contextual decisions relating to how they are used.*” Similarly, we suggest that NIST add the following sentence to the AI Development description: “*Developers provide the initial infrastructure of AI systems, but deployers are responsible for contextual decisions relating to how they are used.*”

BSA would be pleased to serve as a resource for further consultation as you continue to refine the framework.

Sincerely,

Shaundra Watson

Shaundra Watson