BR Business Roundtable

1000 Maine Avenue, NW, STE 500
Washington, DC 20024
202.872.1260
brt.org

January 25, 2022

Ms. Elham Tabassi
National Institute of Standards and Technology
100 Bureau Drive, Stop 200
Gaithersburg, MD 20899

Comments of Business Roundtable
Submitted to AIframework@nist.gov

Re: "Concept Paper for an AI Risk Management Framework"

Dear Ms. Tabassi,

On behalf of the members of Business Roundtable, an association of chief executive officers of leading U.S. companies representing all sectors of the economy, we appreciate the opportunity to respond to NIST's Request for Information (RFI) on its Artificial Intelligence (AI) Risk Management Framework (RMF). As some of the world's largest developers, deployers, and users of AI across many different industries and use cases, Business Roundtable member CEOs have a strong interest in the development of a RMF that mitigates AI risks while supporting AI innovation in order to fully realize AI's benefits for U.S. society and the economy and advance U.S. AI leadership.[1]

Business Roundtable applauds NIST's ongoing collaboration with the private sector and other stakeholders and its efforts to develop a consensus-driven framework for risk management. The direction of NIST's initial AI RMF proposal is commendable, and we have identified many points of alignment between it and the underlying principles and attributes of good AI governance set out in the Business Roundtable Roadmap for Responsible AI and Policy Recommendations, set to be released on January 26, 2022.[2]

These areas of alignment are of critical importance, and we urge NIST to continue to prioritize and feature the following fundamental aspects of AI risk management and governance in subsequent iterations of the RMF and in other avenues of its AI work:

---

[1] A special topic brief on AI published by Business Roundtable that showcases AI's diverse potential, describes U.S. AI leadership and offers a vision for supportive AI policies can be found here: https://system.businessroundtable.org/app/uploads/sites/4/2019/05/ArtificialIntelligence.pdf.

[2] For links, see www.brt.org/ai.

1. AI risk management is a shared responsibility between the private and public sectors and should be multistakeholder in approach (all parts of the organization, affected communities, and customers in mind).
2. This work is grounded in principles of innovation, trustworthiness, transparency, privacy, safety, security, inclusion, and respect for community.
3. Good governance and investing in AI trustworthiness supports innovation and the broader adoption and distribution of technology's benefits.
4. Organizations and government should take a comprehensive approach to AI-related risk management, accounting for both potential negative and positive/beneficial impacts, in both absolute terms and relative to existing technology and human alternatives.
5. Organizations should focus on mitigating high consequence outcomes rather than low consequence outcomes, and consider probabilities of occurrence as appropriate.
6. To support ease of compliance, risk management and governance approaches should align with existing regulatory and risk management frameworks where possible.
7. Effective and consistent risk mitigation requires a culture of accountable technology governance throughout the AI lifecycle.
8. Given the diverse and evolving nature of AI technologies, their use cases, impacts and risks, any AI risk management approach should be voluntary, iterative, and incorporate different perspectives and disciplines. As such, different organizations will encounter and progress with AI governance and risk management in different ways based on their own unique context.

In addition to these areas of clear alignment between NIST's RMF and Business Roundtable's Roadmap for Responsible AI and Policy Recommendations, Business Roundtable has also identified a number of issues and concepts where we believe additional clarification or elaboration would improve the effectiveness and utility of the RMF.

Concepts That Would Benefit from Additional Clarification

1. **Whether the RMF is intended to be used for a broad range of AI use cases, regardless of whether they are likely to create "harms or inequities" for individuals or groups.**

   Specific recommendations include:

   a. Clarify that the RMF is a flexible tool that may be used and adapted by organizations developing, deploying, or using a broad range of AI use cases requiring very different approaches to risk management, including those with greater potential for "harms or inequities" or risks of error or security vulnerabilities, as well as entire categories of AI applications that are not likely to create such harms or risks.

2. **What is meant by: "NIST is agnostic as to whether duties are dispersed and governed throughout an organization or assumed by an individual without any organizational affiliation."**

   Specific recommendations include:

   a. Clarify this statement to understand NIST's position on whether external third parties can take responsibility for an organization's approach to risk management.
   b. Consider, as reflected in Business Roundtable's AI Initiative, that responsibility for AI governance and risk management:
      i. Resides equally inside and outside of the technology departments of any given organization.
      ii. Should include technical and non-technical teams that are stakeholder inclusive and reflect a diversity of experience and expertise.
      iii. Should engage the most senior leadership of any organization, who can direct resources and accountability to the teams responsible for this work.
      iv. May be performed by self-assessments, whether such work is performed internally or against external guidelines or standards.

3. **More fully articulate what is intended by "measurable criteria," as this could be an interesting tool in some contexts. However, measurable criteria and guidelines for AI standards are still nascent and would benefit from further expansion in this draft.**

   Specific recommendations include:

   a. Establish consensus-based standards and opportunities for measurable criteria that are flexible enough to accommodate different industries and use cases and grounded in a multistakeholder standards-setting process.
   b. Clarify that qualitative discussion and scenario planning can also be valuable and/or more feasible when weighing whether to use an AI system, despite having less quantitative measurability.

4. **The RMF focuses resource allocation on engineers and engineering infrastructure, when in practice much of AI governance is process-oriented and non-technical.**

   Specific recommendations include:

   a. Clarify that resources also should be directed to additional technology governance work in operating functions, as well as in strategy, compliance, legal, sustainability, etc.

Concepts That Would Benefit from Further Expansion

1. **The importance of AI risk management being tailored, flexible and use-case specific in order to keep pace with the evolving landscape and avoid being overly restrictive.**

   Specific recommendations include:
   a. Further emphasize impact analysis that accounts for the advantages and improvements gained through AI. Specify that this requires considering the benefits and risks of AI relative to existing human or technology alternatives.
   b. Further distinguish between the diverse nature of potential risks and impacts across different categories of use cases.
   c. Highlight any existing industry regulations that define or categorize "risk."
   d. Further promote engagement with leadership from industry, academia and community organizations to inform AI best practices and enforcement mechanisms, recognizing substantive differences within and across industries.

2. **The importance that organizations' AI governance structures and operations reflect diverse inputs and participation as an effective means for comprehensive and interdisciplinary impact analysis and risk management.**

   Specific recommendations include:
   a. Expressly identify the different types of diversity – professional experience, subject matter expertise and lived experience – that can reduce risk and augment effectiveness among AI design, development and deployment teams.
   b. Provide more specific language around the governance aspects of the RMF, specifically around how to build interdisciplinary governance teams that reflect different aspects and functions from across an organization so that AI governance is embedded broadly and informed by multiple perspectives.

3. **The RMF does not directly address risks of error, drift, security vulnerabilities, lack of interpretability and other AI product outcomes that may adversely impact individuals at a product quality level, regardless of their impact on broader societal outcomes.**

   Specific recommendations include:
   a. Consider adding language around responsible data collection and use:
      i. Monitor and test data collection methodologies and datasets for accuracy and to avoid unfair bias.
      ii. Design data protection safeguards from the outset, collecting and using sensitive data based on demonstrated need.
   b. Consider adding language to account for the system's impact on security and whether it introduces security risks.
   c. Consider adding language around the need to design and implement security safeguards from the outset to protect against adverse manipulation and preserve the system's intended function.

4. **Further illustrate how the concepts of "Core (Functions, Categories and Subcategories)," "Profiles," and "Implementation Tiers" fit within or advance the overall objectives of the initiative.**

   Specific recommendations include:
   a. Provide more details or demonstrations, perhaps with reference to specific use cases, to test or progress this approach.
   b. Consider calling "Categories," within "Functions," "Steps" or "Approaches." This may help create clarity for less technical audiences doing AI governance work.
   c. Discuss that RMF concepts should be aligned or compatible with security and privacy frameworks.

5. **The RMF could usefully be expanded to include earlier technical and non-technical evaluations of AI solutions relative to use case selection and historical performance with existing solutions, and to focus not only on mapping, but also on the alignment of technology goals and broader enterprise goals.**

   Specific recommendations include:
   a. Add to the concept of "Pre-Design," a focus on use case and technology selection and the appropriate alignment of functional goals and model design.

Conclusion

Business Roundtable applauds NIST's proposed AI Risk Management Framework for helping advance Responsible AI efforts. The responsible development and use of AI will realize AI's benefits for society while addressing potential challenges and risks associated with AI design, development and deployment. For this reason, this important issue warrants CEO-level engagement and leadership, as well as participation and engagement from the public and private sectors.

Business Roundtable looks forward to continued engagement with NIST and other thought leaders and policymakers on these important topics.

Sincerely,

William L. Anderson

William L. Anderson
Vice President, Technology and Innovation Policy
Business Roundtable