

February 2, 2024

The Honorable Laurie E. Locascio
Director, National Institute of Standards and Technology
100 Bureau Drive, Mail Stop 8970
Gaithersburg, MD 20899-8970

Re: Business Roundtable Response to Request for Information (RFI) Related to NIST's Assignments Under Sections 4.1, 4.5 and 11 of the Executive Order Concerning Artificial Intelligence

Dear Dr. Locascio,

These comments are submitted on behalf of Business Roundtable, an association of more than 200 chief executive officers (CEOs) of America's leading companies, representing every sector of the U.S. economy. Business Roundtable CEOs lead U.S.-based companies that support one in four American jobs and almost a quarter of U.S. GDP. We appreciate the opportunity to comment in response to the National Institute of Standards and Technology's (NIST) Request for Information (RFI) related to its assignments under Sections 4.1, 4.5 and 11 of the Executive Order Concerning Artificial Intelligence.

Introduction

Business Roundtable member companies across sectors—technology, communications, retail, financial services, health, public safety and security, defense, manufacturing, hospitality, insurance, and others—rely on data and data-driven processes to create, deliver, and improve innovative products and services. Rapid innovation and adoption of AI is transforming the nature of work across every industry and reshaping how people interact with and experience the world around them. AI technologies not only help businesses deliver smarter products and services to their customers but also have enormous potential to drive broader positive change for Americans' health, safety and prosperity.

Our members are among the world's largest developers and deployers of AI. Accordingly, they have a strong interest in ensuring that AI systems are developed responsibly in a manner that centers on principles of innovation, trustworthiness, transparency, privacy, safety, security, inclusion and respect for community.

Business Roundtable encourages policymakers to account for the complex, context-dependent, and rapidly evolving AI ecosystem when developing legislation, regulations, standards, and frameworks. Regarding this RFI and other work around AI, Business Roundtable encourages NIST to:

- Create effective guidelines that are broadly applicable where possible, but that can be tailored for deployment contexts and use cases, or specific types of AI (e.g., generative AI, Machine Learning, Natural Language Processing, etc.) where appropriate;
- Establish a vision for America's role in global AI leadership and define clear goals for how to achieve it while ensuring international standards harmonization as appropriate; and
- Clarify how its work under EO 14110 relates, coordinates and overlaps with the existing work related to AI, including generative AI.

Business Roundtable applauds NIST for soliciting stakeholder comments on how to advance the discussion around responsible development and use of AI. As noted in our responses to the National Telecommunications and Information Administration's (NTIA) Request for Comment on AI Accountability Policies¹, the Office of Science and Technology Policy's (OSTP) Request for Information (RFI) on National Priorities for Artificial Intelligence², and NIST's own RFI on the AI Risk Management Framework (RMF) Concept Paper³, we believe that advancing safe and trustworthy AI is a responsibility shared between all stakeholders, including government and the private sector.

We encourage NIST to take additional ongoing input on these matters given their importance, the broad scope of the feedback requested in this RFI and the complexity of the subjects. While we understand that the deadlines in EO 14110 are ambitious and the issues at hand require urgency, we encourage NIST to complete this work carefully and to take the time that is merited.

Below, we provide comments on NIST's three broad areas of focus in the RFI.

¹ Business Roundtable. (June 12, 2023). Business Roundtable Comment Letter in Response to the NTIA's Request for Comment on AI Accountability Policy. <https://www.businessroundtable.org/business-roundtable-comment-letter-in-response-to-the-ntias-request-for-comment-on-ai-accountability-policy>

² Business Roundtable. (July 7, 2023). Business Roundtable Comment Letter in Response to the Office of Science and Technology Policy's Request for Information on National Priorities for Artificial Intelligence. <https://www.businessroundtable.org/business-roundtable-comment-letter-in-response-to-the-office-of-science-and-technology-policys-request-for-information-on-national-priorities-for-artificial-intelligence>

³ Business Roundtable. (January 25, 2022). Business Roundtable Comment Letter in Response to NIST's Concept Paper for an AI Risk Management Framework <https://s3.amazonaws.com/brt.org/BRT--ResponsetoNISTConceptPaperforanAIRiskManagementFramework.pdf>

Section I: Developing Guidelines, Standards and Best Practices for AI Safety and Security

Well-established standards can empower innovators to work quickly and confidently with existing and emerging technologies. Business Roundtable is supportive of NIST's work developing AI standards and sees a high degree of alignment between it and the *Business Roundtable Policy Recommendations for Responsible Artificial Intelligence*.⁴ Business Roundtable encourages NIST to work to identify and distinguish between use cases and applications, accurately assess their risks, and properly engage key stakeholders in an ongoing discussion around AI governance. NIST must strike a delicate balance between making its guidance overly specific - where it would only aid a select few stakeholders or certain kinds of AI systems - and too vague - where it would serve little real purpose or create uncertainty.

NIST AI RMF Generative AI Companion Resource

As it develops a generative AI companion to the AI RMF, Business Roundtable urges NIST to include several things:

- **Clarify the scope of generative AI guidance.** Generative AI guidance should not depart from the broader context of the AI RMF. Guidance, standards and best practices should consider LLMs and other generative models as different building blocks of AI.
- **Complement existing workstreams and frameworks.** NIST should build upon existing AI risk management practices and standards, as well as related frameworks, and avoid contradicting established guidance, duplication or fragmentation. The generative AI companion should explain how NIST envisions the companion complementing other AI work underway and whether it overlaps with ongoing work of the Generative AI Public Working Group spearheading the development of a cross-sectoral AI RMF profile. Business Roundtable encourages NIST to create resources and cross-walks that reference other frameworks as a generative AI companion to the AI RMF. Broadly, NIST should work to create a consistent and complementary set of AI guidance, standards and best practices so that frameworks do not fragment or become overly difficult to navigate but remain useful for the context at hand.
- **Avoid an overly prescriptive approach.** Some risks of generative AI are well suited to be addressed through broad or general recommendations (e.g., certain types of bias associated with internet-sourced data, drift). But while measures can be taken to protect against these general risks, NIST risk management standards and guidance should avoid a one-size-fits-all approach to the specifics and details of how these items are implemented. For example, for AI models used in a business-to-business context, the

⁴ Business Roundtable. (January 26, 2022). Business Roundtable Policy Recommendations for Responsible Artificial Intelligence.

https://s3.amazonaws.com/brt.org/Business_Roundtable_Artificial_Intelligence_Policy_Recommendations_Jan2022_1.pdf

most effective version of risk management could involve customized risk assessments by the user, in addition to upstream steps taken by developers or deployers.

Implementation should be appreciated as risk- and context-dependent, and flexibility should be afforded to ensure recommendations are effective while allowing innovation.

- **Allow flexibility for implementation and maturity.** Different organizations will encounter and progress with AI governance and risk management in different ways based on their own unique context, which should be reflected in NIST guidance. In addition, NIST should consider the objectives behind the use of a technology when developing guidance. For example, different considerations should be given to AI systems implemented for “defensive purposes” (e.g., cybersecurity, fraud prevention, protection of critical infrastructure). Even targeted guidance addressing specific technologies should be flexible, risk-based and context dependent.

Guidance and Benchmarks for Evaluating and Auditing AI Technologies

Guided by the *Business Roundtable Roadmap for Responsible Artificial Intelligence*⁵, our members are already working to refine their existing internal risk management strategies by establishing cross-functional AI ethics and governance committees, conducting regular fairness and ethics assessments, and implementing performance monitoring processes. Business Roundtable members see alignment between this work and NIST’s efforts to create guidance and benchmarks for evaluating and auditing AI models.

To ensure that these guidance and benchmark documents are successful, Business Roundtable encourages NIST to:

- **Develop a nuanced concept of risk.** Complex models and complex use cases will require nuanced risk management throughout the lifecycle of AI models and systems. The discussion of risk in EO 14110 does not appropriately capture the need for this context-aware nuance. For example, EO 14110’s computational thresholds for reporting on dual-use systems, which are nominally based on their risk to national security, do not reflect true risk. Small, narrowly capable models can carry much more risk than broad ones depending on their application and can require significantly less data, computing power, and time to train. As NIST aims to develop a nuanced concept of AI risks, it should also compare the risks, impacts and outcomes of AI to existing human or technological alternatives.
- **Offer options for transparency, including for different audiences.** The use of model cards, data cards, system cards, benchmarks, impact assessments, and other forms of documentation demonstrate the variety of transparency tools that may be used to understand AI system behavior, assess the potential for bias, and ensure accountability.

⁵ Business Roundtable. (January 26, 2022). Business Roundtable Roadmap for Responsible Artificial Intelligence. https://s3.amazonaws.com/brt.org/Business_Roundtable_Artificial_Intelligence_Roadmap_Jan2022_1.pdf

Various kinds of transparency are crucial for building trust with regulators, end users, and other members of the public, and these tools may be tailored to the needs and interests of these different audiences. However, NIST should also carefully consider how to strike a balance between transparency necessary to foster public trust, safeguarding intellectual property essential to drive future innovation and the risk of disclosing information that could inadvertently aid bad actors.

- **Work collaboratively to create guidance on how to address unfair and harmful bias in AI development and deployment.** Establishing processes to mitigate, detect and remediate unfair or harmful algorithmic and AI bias, particularly in areas lacking clear best practices and established conventions (e.g., methodologies and frameworks for identifying and mitigating proxy bias) is essential. Adding specificity and clarity regarding operationalization of the socio-technical aspects of the RMF would help companies calibrate their compliance efforts and encourage innovative approaches to mitigating unfair and harmful bias. Established bias reduction methodologies (disparate impact assessment, equal opportunity difference, etc.), technical guides, and research papers exist on how to effectively combat bias in sociotechnical systems. NIST should seek additional input from deployers, developers and other experts, while leveraging the existing body of work wherever possible. When drafting guidance, NIST should focus on filling in gaps where clear guidance does not already exist.
- **Emphasize the importance of a diversity of skills, background and expertise.** Business Roundtable members rely on countless individuals with diverse professional experience, disciplinary expertise and backgrounds to govern AI. NIST AI guidance documents should emphasize the importance of this diversity on these interdisciplinary teams implementing AI governance. NIST should also consider laying out a high-level approach to training and equipping individuals from diverse professional backgrounds to constructively and effectively participate in AI governance.

Guidelines for AI Red-Teaming Tests

Within the EO, the term “red-teaming” is quite broad, including “structured testing effort to find flaws and vulnerabilities in an AI system” including “harmful or discriminatory outputs from an AI system, unforeseen or undesirable system behaviors, limitations, or potential risks associated with the misuse of the system.” This encompasses a number of distinct kinds of evaluation, and NIST should clarify any differences between traditional definitions of red-teaming for cybersecurity and the EO’s definition. In addition, Business Roundtable encourages NIST to:

- **Clearly define goals of red-teaming.** The RFI mentions a selection of the risks that red-teaming tests might evaluate, but red-team tests as defined in the EO are so broad that any testing to find flaws or vulnerabilities would qualify. NIST should focus on developing specific testing guidance that targets efforts on highest priority areas of

concerns, while acknowledging that the red-teaming process will need to be adapted to different sectors, types of AI models and use case contexts.

- **Focus on concrete risks.** Like other risk evaluation, AI red-teaming exercises, and the risks they assess, and guidance for this evaluation should be focused on immediate and concrete risks, and the goal of these evaluations and tests needs to be clear. The goal of enabling developers to find flaws and vulnerabilities in AI systems is not well served by such a broad set of criteria and evaluation.
- **Place red-teaming in the context of holistic risk management.** Red-teaming is an important aspect of holistic AI risk management, but certain risks (e.g., data and methodology-related risks) also require other approaches. NIST's guidance should reflect these facts and avoid an over-emphasis on red-teaming in the context of comprehensive risk management. It should also develop guidance in a spirit of innovation and avoid prescribing an inflexible check-the-box exercise (e.g., laying out what red-teaming might look like at different stages of the development pipeline).

Additional Considerations

Business Roundtable member companies across sectors face significant and growing cybersecurity threats, including in their use of AI technologies and systems. NIST should ensure AI-focused cybersecurity guidance is consistent with existing cybersecurity guidance and standards, as well as the agency's ongoing work with CISA.

In addition, NIST's guidance should also contemplate other societal and individual risks to operating safe, secure and trustworthy AI systems. Acknowledging the broad diversity of risk rather than focusing on a few risks within silos will be critical. These risks should be within scope of the generative AI companion and other guidance relating to evaluation and auditing capabilities for AI technologies.

Section II: Reducing the Risk of Synthetic Content

The current risk landscape of synthetic content is dynamic and complex. AI systems, and synthetic content in particular, amplifies existing problems that already occur across all aspects of society (i.e., misinformation, market disruptions, human rights problems, etc.). Business Roundtable supports efforts to understand and reduce the risk of content that has been created or modified through the use of AI. NIST should clearly define the scope of synthetic content and the risks it aims to address and ensure that approaches are well tailored to these risks.

As NIST develops these standards, Business Roundtable encourages NIST to consider:

- **Synthetic content is not necessarily a risk.** Detection of synthetic content does not necessarily identify whether that content poses a risk. In many cases, the risk of the

content may not be based on the synthetic nature of the content but how it is being used. For example, there are uses for synthetic content in customer service systems and other applications that are trustworthy. Similarly, many public awareness campaigns rely on compelling and credible synthetic content, both generated by AI and otherwise modified by AI.

- **Risk can be addressed in many ways.** Risks associated with AI-created content can be addressed when a model is developed, as part of the interface created when the model is deployed, or as output content is used and distributed. NIST should clarify that it will focus on the most effective ways to address the most common and impactful risks. An approach that focuses too heavily on ensuring that models do not create harmful content may stifle innovation, even though that content could be mitigated at deployment. Additionally, any risk mitigation that relies on models or content creators to label synthetic content will fail where bad actors work within the system.
- **Provenance and authenticity may be more important than how content was created.** Being able to determine where content is from (its provenance) and whether a given person or organization shared it (its authenticity) may be more relevant to evaluating or mitigating risk of that content than determining whether content is “synthetic.” In a world where the average internet user has access to advanced generative AI tools, synthetic content is already ubiquitous and used for both positive and negative aims. Initiatives focused on provenance and authenticity can help ensure that people understand who shared content and whether they should trust it.

Section III: Advance Responsible Global Technical Standards for AI Development

Business Roundtable strongly supports the development of international standards for AI and has called for global coordination on key issues in our 2022 Policy Recommendations for Responsible Artificial Intelligence.⁶ Three of our ten policy recommendations focus on international standards and collaboration, and we encourage NIST to keep them in mind:

- Prioritize strategic international engagement on AI issues;
- Engage on global AI standards and guidelines; and
- Strive for common principles and interoperability.

International alignment among U.S. allies on consensus standards when possible will not only help safeguard against risks but also spur technological development. When researchers have a clear regulatory environment with established frameworks, they can more confidently pursue

⁶ Business Roundtable. (January 26, 2022). Business Roundtable Policy Recommendations for Responsible Artificial Intelligence.

https://s3.amazonaws.com/brt.org/Business_Roundtable_Artificial_Intelligence_Policy_Recommendations_Jan2022_1.pdf

their research objectives. Moreover, investors can feel more comfortable funding projects without fear that they will be shut down later.

To support the responsible development of AI technologies, Business Roundtable encourages NIST to:

- **Align key definitions.** Currently, developers and policymakers suffer from a lack of standardized definitions in the field of AI. Key terms such as explainability, classification of risk, developer and deployer, and even “AI” itself are used in different ways by different actors. By pushing for alignment between the United States and its allies on key definitions when possible, NIST can provide a solid foundation for researchers and ensure that future policy debates are properly scoped and drive towards pragmatic solutions.
- **Adopt flexible and aligned governance approaches.** Even among friends and allies, regulatory approaches to AI will not be exactly the same. Policymakers in different jurisdictions will rightly center the concerns that matter most to their constituencies, and standardizing regulatory and governance approaches across all jurisdictions is impractical. NIST should focus on promoting interoperability, flexibility and minimizing direct conflicts when appropriate among different regulations through global standards-setting bodies and regulatory forums, working closely with private-sector stakeholders.
- **Recognize the importance of cross-border data flows and digital trade frameworks.** In recent years, many governments have implemented data localization requirements that force businesses to store data on servers within their jurisdiction. In addition to lowering productivity and raising prices, this trend also has particularly negative consequences on AI. By impairing cross-border data flows, data localization prevents AI researchers from accessing diverse and representative data that they could use to train their models. To ensure that future AI models are fair and representative of all groups, the openness of cross-border data flows wherever possible is essential.

Conclusion

Business Roundtable welcomes NIST’s work to identify the technical and policy efforts needed to cultivate trust in AI. Business Roundtable looks forward to continued engagement with NIST and other thought leaders and policymakers on these important topics and is happy to discuss our response or these issues at any time. Please contact Amy Shuart, Vice President of Technology & Innovation, Business Roundtable, at ashuart@brt.org or (202) 496-3290.