U.S. Chamber of Commerce
Technology Engagement Center

MICHAEL RICHARDS
Director

1615 H STREET, NW
WASHINGTON, DC 20062-2000
(202) 463-5518
MRichards@uschamber.com

September 15, 2021

National Institute of Standards and Technology
U.S. Department of Commerce
100 Bureau Drive, Stop 2000
Gaithersburg, MD 20899

Re: Artificial Intelligence Risk Management Framework Request for Information

To Whom It May Concern:

The U.S. Chamber of Commerce's Technology Engagement Center ("C_TEC") appreciates the opportunity to submit feedback to the National Institute of Standards and Technology ("NIST") in response to its Request for Information (RFI) about its development of an "Artificial Intelligence Risk Management Framework" ("AI RMF").

C_TEC commends NIST for taking the lead in bringing together stakeholders to "help inform, refine, and guide the development of the AI RMF" to assist in the trustworthy and responsible use of AI technologies[1]. C_TEC strongly agrees that a stakeholder-led, consensus-based framework can promote digital innovation within the AI field. This is further indicated in C_TEC's recently published report on trustworthy AI[2], where 54% of respondents indicated that they support open source tools and frameworks facilitated by the government can help enable the development of new AI technologies. There is a broad understanding that a balanced and innovative framework can help mitigate any risks posed by AI. This is why C_TEC has continued to advocate for the establishment of an AI Risk Management Framework. [3]

C_TEC recognizes NIST's ongoing engagement across stakeholders and the consensus-driven processes leveraged to develop the voluntary framework as well as other NIST work around AI principles. C_TEC looks forward to upcoming opportunities, such as workshops, to participate and collaborate with NIST and cross-industry stakeholders.

C_TEC appreciates the opportunity to comment on NIST's request for information to develop a voluntary AI RMF. C_TEC offers the following comment to the corresponding questions within the RFI for the development of the voluntary AI RMF.

---

[1] https://www.federalregister.gov/documents/2021/07/29/2021-16176/artificial-intelligence-risk-management-framework

[2] https://americaninnovators.com/wp-content/uploads/2021/07/AIInstitute_ChamberofCommerceReport_v27-pages.pdf

[3] https://americaninnovators.com/wp-content/uploads/2020/10/CTEC_TechUpgrade_Data_.pdf

1. The greatest challenges in improving how AI actors manage AI-related risks—where "manage" means identify, assess, prioritize, respond to, or communicate those risks;

C_TEC understands that there is a wide array of issues related to managing AI-related risk. Some challenges result from resource constraints and incentives that make investing in AI risk management difficult, while the rapid pace of technological evolution has made it difficult for skills, regulations, and professional credentialing to keep up. One of the most significant challenges will be improving the management of AI-related risk due to its regulatory infancy.

Furthermore, we also believe that there is a need for further technical expertise within federal regulating agencies. This is why C_TEC is pleased to see the AI in Governance Act was signed into law. This vital piece of legislation will help identify the current needs within the federal workforce, which will then assist in helping agencies hire the necessary expertise to help improve how AI actors can manage AI-related risks.

C_TEC believes that a growing patchwork of local, state, federal, and international regulations and standards around AI challenges AI actors and their ability to manage AI risks effectively. An overcomplicated regulatory environment potentially inhibits effective management. International standards serve as a tool for harmonizing regulatory processes to ensure greater interoperability and avoid a fragmented global network of differing regulations. This is why C_TEC supports efforts to abide by international standards[4].

Additionally, NIST should consider exploring the inherent contradictions and subsequent tradeoffs between their characteristics when developing the AI Risk Management Framework. For example, principles such as privacy and explainability can be contradictory when personal data is involved. An overemphasis on the explainability of an AI system can lead to diminishing privacy. If explainability regulations required detailed information regarding the training dataset being used, the type of machine learning algorithm, or other information, the regulations could make attacks on machine learning models, such as model extraction, much easier and more successful. In the absence of such details, attackers are left to less successful black-box attacks. How organizations weigh these various tradeoffs should be a subject of discussion for NIST as it develops the framework.

3. How organizations currently define and manage principles of AI trustworthiness and whether there are important principles which should be considered in the Framework besides;

In this request for information, NIST appears to be making a distinction between so-called "characteristics" and "principles" of AI trustworthiness. C_TEC believes that NIST should provide more details on this distinction in future requests and workshops. For example, some may see "privacy" – currently classified as a characteristic – as fundamental to successfully developing and deploying AI. C_TEC recommends NIST further clarify if "privacy" is considered a principle or characteristic.

---

[4] https://www.uschamber.com/press-release/us-chamber-releases-artificial-intelligence-principles

4. The extent to which AI risks are incorporated into different organizations' overarching enterprise risk management—including, but not limited to, the management of risks related to cybersecurity, privacy, and safety;

In considering how to incorporate AI risks into organizations' enterprise risk management, NIST should carefully consider the scope of AI risk. There are risks that are of particular concern with regards to AI, such as harmful bias and model drift. However, many risks that may be associated with AI may come more generally from existing issues being propagated by digital transformation and not be intrinsically tied to AI. NIST should consider how more general risks and AI-specific risks are addressed in the Risk Management Framework.

C_TEC believes that there are opportunities to integrate AI risk management into existing processes where they already exist. For example, well-developed risk and compliance processes already exist within the financial services industry, and AI risk management should be integrated into these pre-existing structures. As AI is embedded into every part of the business, lines of business using AI need to customize and adopt their own risk management but base guidance from a central team that may be part of security, legal, responsible business, and data and AI center of excellence.

However, many organizations in fields with less robust risk and compliance processes will need to create new risk management structures and processes. To best assist organizations that may find themselves building out new risk and compliance processes for AI risk, NIST should provide examples of frameworks that can be used to develop and implement AI risk management processes. Further, examples of different maturity models for AI risk management may be helpful (e.g., per industry or business function).

5. Standards, frameworks, models, methodologies, tools, guidelines and best practices, and principles to identify, assess, prioritize, mitigate, or communicate AI risk and whether any currently meet the minimum attributes described above;

C_TEC firmly believes that safe, ethical, and effective AI systems can provide tremendous benefits to society. However, we understand that AI is increasingly being developed and deployed within critical processes (e.g., healthcare, employment, judicial, policing, etc.) where there is a concern that such systems could pose a risk to safety, privacy, and human rights. Therefore, it is important to evaluate the level of risk posed by AI and its intended application to determine an appropriate course of action for mitigating any existing or potential risks. C_TEC recommends conducting a "risk triage" to determine whether the risk posed by a particular use of AI is low or high, as an organization will need to put in place different processes to deal with low- and high-level risk.

Moreover, as AI is developed across data cooperatives where multiple lines of business or even partners work together to supply data and create algorithms, standards that include common measures and metrics are essential to enable efficient operations.

7. AI risk management standards, frameworks, models, methodologies, tools, guidelines and best practices, principles, and practices which NIST should consider to ensure that the AI RMF aligns with and supports other efforts;

While countries around the world will undoubtedly pursue their own ways of regulating AI, NIST should strive to harmonize definitions of key terms with those already published so that the global AI community is speaking the same language. A common lexicon will give organizations and society more confidence and promote greater alignment of standards, frameworks, models, etc.

Legislative bodies, regulatory agencies, standards bodies, and others have started developing their own definitions related to AI, the proliferation of which unnecessarily complicates operations for global organizations. For example, Article 3 of the European Commission's proposed Artificial Intelligence Act includes numerous definitions of AI-related terms, including 'artificial intelligence system' and 'training data.' In the United Kingdom, the Information Commissioner's Office has published its own definition of artificial intelligence. In the United States, Congress defined artificial intelligence in the John S. McCain National Defense Authorization Act. The Institute of Electrical and Electronics Engineers (IEEE) has also weighed in with its own definitions. Ensuring some degree of harmony in the terms used to talk about AI is one way that NIST's Risk Management Framework can promote innovation and adoption of AI.

Additionally, NIST should note the increasing number of countries that are incorporating the concept of risk triaging into their frameworks, including both Canada and the European Commission. By distinguishing between high-risk and low-risk uses of AI, regulators are able to focus their limited resources on the uses of AI that could have the most significant impacts on individuals' lives.

9. The appropriateness of the attributes NIST has developed for the AI Risk Management Framework.

C_TEC applauds and aligns with the attributes for the AI RMF NIST has outlined, and offers the following considerations or additions to the attributes as NIST develops the framework:

Consensus-driven and developed, and regularly updated through an open, transparent process.

C_TEC supports developing a consensus-driven AI RMF through an open, transparent process and believes it is a highly effective means of addressing the challenges and opportunities presented by emerging technologies such as AI.[5] NIST should build on its previous experience in developing the successful Cybersecurity Framework and replicate in the development of the Privacy Framework, where a consensus-driven, open, and transparent process was used to create and maintain consumer and stakeholder trust throughout and after the process. We encourage NIST to align the RMP with these frameworks to ensure consistency across AI, privacy, and cybersecurity, which have significant overlap in issues and governance processes. Furthermore,

[5] https://americaninnovators.com/news/u-s-chamber-releases-artificial-intelligence-principles/

C_TEC notes that multi-stakeholder initiatives have the most considerable capacity to enable industry harmonization and provide equal access to AI resources across stakeholders.

One mechanism for such a multi-stakeholder initiative is policy prototyping through which different stakeholders can co-create and test the AI RMF. Policy prototyping is an experimentation-based approach for policy development that can provide a safe testing ground to test and learn early in the process how different approaches to the formulation of the AI RMF might play out when implemented in practice while assessing their impact before its actual release. Policy prototyping involves a variety of stakeholders that come together to co-create governance frameworks, including regulation and standards. Developing and testing governance frameworks in a collaborative fashion, this allows policymakers to see how such frameworks can integrate with other co-regulatory tools such as corporate ethical frameworks, voluntary standards, certification programs, ethical codes of conduct, and best practices, as referenced below. This method has been successfully been used in Europe to test an AI Risk Assessment framework, leading to several concrete recommendations for improving self-assessments of AI.

Provide common definitions

C_TEC supports the development of common definitions. AI systems are being used in a wide array of applications from a diverse group of stakeholders, which is why it is important to develop common definitions among private and public sector stakeholders. One major impetus for establishing common definitions is that there is a significant risk that agencies may adopt different definitions relating to AI, leading to unnecessary confusion. With several agencies as well as State Agencies and legislatures currently in the process of reviewing and developing their own governance approaches to AI, C_TEC recommends that NIST make a concerted effort to coordinate with other federal government agencies to facilitate collaboration to help harmonize common definitions on AI.

When addressing definitions, C_TEC urges NIST to avoid defining premature requirements relating to aspirational standards for AI that either have not matured or are not developed yet, despite ongoing efforts by the private sector and academia to develop them. These include concepts such as "explainability," "auditability," "robust accuracy," and "error-free algorithms." If the RMF defines and deploys these concepts beyond their current development, as it will add unnecessary burdens which could stifle innovation. In the immediate term, the focus should be tied to specific and concrete harms, like unlawful discrimination. We would also express the need for NIST to continue to update terms as technology continues to evolve.

Use plain language that is understandable by a broad audience

We are encouraged to see the emphasis on developing an AI RMF in an understandable and comprehensible manner. As AI becomes further present in our daily lives, any framework must be developed using terms and vocabulary that is understandable and will allow for the American public and other stakeholders to engage meaningfully and provide necessary feedback. Furthermore, we believe this attribute is an essential part of establishing public trust in the framework.

Be adaptable to many different organizations

AI is a changing field in which new, transformative technologies are regularly emerging. Moreover, AI is a tool that has a multitude of diverse applications, making a one-size-fits-all and prescriptive framework particularly ill-suited for AI. Many promising uses of AI relate to the improvement of business processes and operations, which present a low risk of harm to individuals. This was recently recognized in the European Commission's draft Artificial Intelligence Act, which, despite other flaws, appropriately recognized that many AI applications present a low risk that should benefit from a light regulatory touch. C_TEC supports a flexible, non-prescriptive framework that can adapt to many different organizations and industries. Furthermore, any framework should be able to be applied broadly and be scalable, given the diversity of the size of businesses and types of sectors utilizing and developing AI.

Be risk-based, outcome-focused, voluntary, and non-prescriptive.

C_TEC strongly supports a risk-based framework as detailed in the U.S. Chamber's Artificial Intelligence Policy Principles[6]. We believe that any risk-based approach should account for the varying magnitude and nature of consequences when considering risk mitigation, recognize tradeoffs (including explicit consideration of benefits as well as risks in designing risk management approaches), and ensure that any regulatory approaches being considered or proposed are linked to specific public policies in the national interest.

We urge NIST to focus on risk mitigation rather than the elimination of risk as contemplated in the RFI where it defines "responding" to include "avoiding risk." Avoiding risk is an infeasible metric for any technology, system, or program, including AI. In the context of AI specifically, it is impracticable to ensure that AI outcomes can be accurate and robust even where there are *no* errors in data sets. Imposing an infeasible metric like "avoiding risk" or being "free of errors" will stymie AI innovation. The RMF should focus instead on "responses" that aim to mitigate risks, and the reasonable design, oversight, and monitoring that can achieve robust mitigation.

Furthermore, C_TEC supports the development and implementation of a framework that is voluntary. A voluntary consensus framework can help create and safeguard trust at the heart of AI-driven systems and business models and permit the flexibility for innovation, allowing the framework to develop with the technology.

The RMF should, furthermore, explicitly acknowledge that not all AI risks can be effectively identified or measured and should not restrain AI innovation as a result. Standards, guidelines, and best practices in this emerging technology area are still in the process of being developed. Because of this, we are still learning about the range of potential risks, their likelihood, and how to measure them. The RMF should specifically address situations where risk cannot be measured and offer guidance on reasonable steps for mitigating that risk without limiting innovation and investments in new and potentially beneficial AI technologies.

Be readily usable as part of any enterprise's broader risk management strategy and processes.

[6] U.S. Chamber Releases Artificial Intelligence Principles | U.S. Chamber of Commerce (uschamber.com)

C_TEC believes that any opportunity to limit redundancy or unneeded processes for businesses should be encouraged. Many businesses already have risk management processes and procedures in place for AI. This is why we believe it's vital that NIST asks for stakeholder input on what current risk management processes are being used and common methods that could be used more broadly.

Be consistent, to the extent possible, with other approaches to managing AI risk.

C_TEC agrees that any framework should provide for consistency as applications vary across industries and the federal government. Many different stakeholders have previously worked closely with their regulators to develop processes to manage AI risk. This is why we believe it is important that NIST works closely with those agencies to be consistent as possible with other policies and procedures that are in place.

For example, the financial services industry is already heavily regulated and has existing risk management frameworks in place to manage risks associated with AI. The AI RMF should recognize this and not seek to apply duplicative standards to an already heavily regulated sector. With this in mind, we refer NIST to SR-11-7 to consider leveraging such model risk management principles, e.g. conceptual soundness for AI developed outside of the financial services industry.

Be a living document.

C_TEC understands that AI is an everchanging field and that emerging technologies are continuously evolving, requiring any framework to be routinely reviewed and updated. However, any review must be conducted with stakeholder input and be transparent to ensure continued trust in the framework. NIST should also seek public comment on whether the AI RMF needs to be updated before pursuing revisions.

Furthermore, we would encourage NIST to continue to work closely along with stakeholders to determine how the input formulation of the AI RMF might play out when implemented in practice while assessing their impact before its actual release.

11. How the Framework could be developed to advance the recruitment, hiring, development, and retention of a knowledgeable and skilled workforce necessary to perform AI-related functions within organizations.

Scaling AI effectively for the long term will require the professionalization of the industry, and NIST can help advance this goal by including recommendations and best practices for the professionalization of organizations' AI workforces in the AI RMF. Stakeholders – from practitioners to leaders across the private and public sector – must come together to distinguish clear roles and responsibilities for AI practitioners; demand the right level of education and training for practitioners; define processes for developing, deploying and managing AI; and democratize AI literacy across the enterprise. By formalizing AI as a trade with a shared set of norms and principles, companies will be poised to achieve more value from AI.

Real value can only be realized when trained AI practitioners are working hand in hand with the business to accomplish their organization's goals, and those interdisciplinary teams are guided by standards, rules, and processes. By following these steps to standardize professionals and processes, organizations can better set themselves up to scale AI and, in so doing, make the most of this quickly evolving technology.

1. **Distinguish clear AI roles**: A hallmark of a professionalized industry or trade is that practitioners understand the individual roles that contribute to a final product. Multidisciplinary teams of diverse perspectives, skills, and approaches must work together to innovate and deliver AI products or services. The mix and the ratio of roles is going to depend on the use cases pursued at the time and will vary from project to project. Establishing a blueprint for how teams should operate will help this process become more turnkey over time. Yet one thing remains true across all projects – organizations need to establish ownership and expectations from the start.

2. **Define AI processes**: While some argue that formalized processes and governance could stifle innovation, research has shown the opposite. In professionalized industries, there's a standard approach to testing and benchmarking during the creation (or optimization) of products and services. Similarly, whether a company is making smart devices or building a data science model to improve the online retail experience, establishing systems and processes to support the development of the AI product or solution allows people to innovate in a predictable and efficient way.

Conclusion

C_TEC appreciates NIST's ongoing efforts to improve the management of risk to individuals, organizations, and those associated with AI by creating a voluntary Risk Management Framework. Establishing a voluntary Risk Management Framework has significant promise in creating an innovative environment for Artificial Intelligence, which is why we urge NIST to continue to work closely with stakeholders to ensure that innovation is not stifled. We thank you for your consideration of these comments and would be happy to further discuss any of these topics.

Sincerely,

Michael Richards
Director, Policy
Chamber Technology Engagement Center