**Before the Department of Commerce**
**National Institute of Standards and Technology**
**Washington, D.C.**

In the Matter of

| | | |
|---|---|---|
| NIST AI Risk Management Framework | ) | Initial Draft |
| | ) | |

# COMMENTS OF CTIA

Thomas K. Sawanobori
Senior Vice President and Chief Technology Officer

John A. Marinho
Vice President, Technology and Cybersecurity

Melanie K. Tiano
Assistant Vice President, Cybersecurity and Privacy

Avonne S. Bell
Director, Connected Life

**CTIA**
1400 16th Street, NW, Suite 600
Washington, DC 20036
202-736-3200
www.ctia.org

April 29, 2022

## Table of Contents

## I. INTRODUCTION AND SUMMARY.

CTIA[1] welcomes the opportunity to provide feedback to NIST regarding the initial draft ("Draft") of the Artificial Intelligence ("AI") Risk Management Framework ("RMF").[2] CTIA has been actively engaged with NIST throughout the AI RMF development process—submitting comments in response to NIST's AI RMF Request for Information and AI RMF Concept Paper, participating in the October 2021 Kickoff Workshop, and facilitating a meeting between NIST's AI RMF development team and CTIA members.[3] We look forward to continued collaboration with NIST.

With these comments, we encourage NIST to focus on creating a risk management tool tailored for AI that builds on the foundation of NIST's prior guidance documents and that is targeted to help organizations practically manage AI implementation. Specifically, CTIA has the following suggestions: (1) NIST should continue to develop the AI RMF as a voluntary, flexible, and risk-based tool, which will best promote innovation in AI; (2) NIST should devote more attention to the benefits of AI; (3) the AI RMF should be a process-based and policy-

[1] CTIA – The Wireless Association® (www.ctia.org) represents the U.S. wireless communications industry and the companies throughout the mobile ecosystem that enable Americans to lead a 21st-century connected life. The association's members include wireless carriers, device manufacturers, suppliers as well as apps and content companies. CTIA vigorously advocates at all levels of government for policies that foster continued wireless innovation and investment. The association also coordinates the industry's voluntary best practices, hosts educational events that promote the wireless industry, and co-produces the Industry's leading wireless tradeshow. CTIA was founded in 1984 and is based in Washington, D.C.

[2] NIST, AI Risk Management Framework: Initial Draft (Mar. 17, 2022), https://www.nist.gov/system/files/documents/2022/03/17/AI-RMF-1stdraft.pdf ("Draft").

[3] Artificial Intelligence Risk Management Framework Request for Information, Department of Commerce, NIST, 86 Fed. Reg. 40810 (July 29, 2021), https://www.govinfo.gov/content/pkg/FR-2021-07-29/pdf/2021-16176.pdf; NIST, AI Risk Management Framework Concept Paper (Dec. 13, 2021), https://www.nist.gov/system/files/documents/2021/12/14/AI%20RMF%20Concept%20Paper_13Dec2021_posted.pdf; Comments of CTIA, NIST, Artificial Intelligence Risk Management Framework Request for Information (Sept. 15, 2021); Comments of CTIA, NIST, AI Risk Management Framework Concept Paper (Jan. 25, 2022) ("CTIA AI RMF Concept Paper Comments").

neutral tool and given that specific characteristics of trustworthy AI are policy-based and nuanced, the AI RMF should establish a flexible approach that can be adapted to different characteristics as appropriate; (4) NIST should take steps to maximize the utility of the AI RFI, such as honing its intended audience, mapping the framework to existing private and public sector resources, and avoiding prescriptive guidance; and (5) NIST should finalize the first version of the AI RMF before building profiles, attempting to measure effectiveness, or otherwise taking next steps that will build upon the AI RMF.

## II. NIST'S DEVELOPING VOLUNTARY, FLEXIBLE, AND RISK-BASED FRAMEWORK IS THE RIGHT APPROACH TO MANAGE RISKS AND MAXIMIZE BENEFITS ASSOCIATED WITH AI.

The AI RMF promises to be a helpful tool for applying longstanding risk management principles and approaches to emerging AI technologies, and the Draft is currently on the right track. The Draft establishes a voluntary approach to AI risk management.[4] It is also flexible and risk-based, acknowledging that there is no one-size-fits-all solution to addressing risk. As the Draft states, "the AI RMF is neither a checklist nor should be used in any way to certify an AI system."[5] The Draft is also intended to be used by organizations of different sizes and across a wide spectrum of AI use cases.[6]

[4] Draft at 1 ("This voluntary framework provides a flexible, structured, and measurable process to address AI risks throughout the AI lifecycle, offering guidance for the development and use of trustworthy and responsible AI."); *see also id.* at 3 (listing as an attribute of the AI RMF that it strives to "[b]e risk-based, resource efficient, and voluntary").

[5] *Id.* at 1. Flexibility is enshrined in two of the main attributes of the AI RMF, which strives to "[b]e easily usable and mesh with other aspects of risk management. Use of the Framework should be intuitive and readily adaptable as part of an organization's broader risk management strategy and processes. It should be consistent or aligned with other approaches to managing AI risks." The AI RMF also aims to "[b]e useful to a wide range of perspectives, sectors, and technology domains. The AI RMF should be both technology agnostic and applicable to context-specific use cases." *Id.* at 3.

[6] *Id.* at 2 ("The NIST AI RMF offers a process for managing risks related to AI systems across a wide spectrum of types, applications, and maturity. This framework is organized and intended to be understood and used by individuals and organizations, regardless of sector, size, or level of familiarity with a specific type of technology. Ultimately, it will be offered in multiple formats, including online versions, to provide maximum flexibility.").

This voluntary, flexible, and risk-based approach will best facilitate innovation and beneficial uses of AI by providing organizations with a framework to manage risks while implementing AI. Indeed, these attributes have been critical to successful NIST guidance such as the Cybersecurity Framework ("CSF") and Privacy Framework. NIST should continue to follow the same tried-and-trusted path with the AI RMF. As the Draft is finalized, NIST should continue to prioritize these attributes, including with respect to forthcoming parts such as the Practice Guide and profiles. AI technology is developing rapidly with a wide range of varied use cases. A voluntary, flexible, and risk-based tool will create an environment for technology development where system risks are considered while innovation is still encouraged.

This approach is not only good policy, but it is also required under the National AI Initiative Act of 2020, which specifically directs NIST to develop the AI RMF. Congress recognized the importance of voluntariness and flexibility, codifying these virtues in the statute. The National AI Initiative Act charges NIST with research and development of "best practices and voluntary standards" for AI systems[7] and explicitly mandates that the AI RMF "not prescribe or otherwise require the use of specific information or communications technology products or services."[8]

Accordingly, NIST is rightly committed to a risk-based, voluntary, and flexible framework,[9] and should continue to prioritize these longstanding principles to enhance the effectiveness of the framework.

[7] William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116-283, § 5301(b)(1), 83 Stat. 4536–4539 (2021) ("2021 NDAA").

[8] *Id.* § 5301(c)(6).

[9] Draft at 3.

## III. THE DRAFT STARTS THE CONVERSATION ABOUT THE POSITIVE IMPACTS OF AI, BUT NIST SHOULD DEVOTE MORE ATTENTION TO UNDERSTANDING AND ASSESSING THE FULL RANGE OF AI BENEFITS.

### A. The Draft Acknowledges that AI Technology Has a Multitude of Benefits.

As CTIA has highlighted throughout this development process, AI systems and technologies promise enormous benefits. Notably, AI promises many benefits for telecommunications and, as the January 2021 FCC report from the industry-led Technological Advisory Council ("TAC") illustrates, AI has broad applicability across the industry.[10] In the wireless sector in particular, AI can be implemented in many use cases, including real-time network threat detection, combatting fraud, customer service, and helping to build 5G networks.[11]

In the Draft, NIST has included several examples of the benefits and potential of AI technology. NIST explains:

> Remarkable surges in artificial intelligence (AI) capabilities have led to a wide range of innovations with the potential to benefit nearly all aspects of our society and economy – everything from commerce and healthcare to transportation and cybersecurity. AI systems are used for tasks such as informing and advising

[10] *See* Technological Advisory Council, The Importance of Artificial Intelligence and Data for the Telecommunications Industry and the FCC, at 5 (Jan. 14, 2021), https://www.fcc.gov/sites/default/files/fcc_aiwg_2020_whitepaper_final.pdf.

[11] AT&T, AT&T Innovations: Raw Data to Real Solutions: How AT&T Applies AI (June 2019), https://policyforum.att.com/wp-content/uploads/2020/08/190531_ATT_Innovations_InfluxDataFueling_05.pdf; Ericsson, *T-Mobile, improving customer experience with AI and IT Operations*, https://www.ericsson.com/en/cases/2021/tmobile-improve-customer-experience-with-ai, (last visited Apr. 26, 2022); AWS, *At T-Mobile, AI Humanizes Customer Service: Using Technology to Improve Personal Connections*, https://aws.amazon.com/machine-learning/customers/innovators/t_mobile/, (last visited Apr. 26, 2022); Kyle Ragonese, *IBM and Verizon Business to collaborate on 5G and AI solutions at the Enterprise Edge*, Verizon (July 16, 2020), https://www.verizon.com/about/news/ibm-and-verizon-business-collaborate; Karen Schultz, *Verizon and Cellwize speed deployment of Verizon's 5G network, simplify development for the network*, Verizon (July 15, 2020), https://www.verizon.com/about/news/verizon-cellwize-speed-deployment; Sara Castellanos, *Verizon Enlists AI in 5G Network Build-out*, Wall St. J. (Aug. 4, 2021), https://www.wsj.com/articles/verizon-enlists-ai-in-5g-network-build-out-11628103712; Akash Palkhiwala, *Qualcomm power-efficient AI: Making Technology more sustainable*, Qualcomm (Nov. 10, 2021), https://www.qualcomm.com/news/onq/2021/11/10/qualcomm-power-efficient-ai-making-technology-more-sustainable.

> people and taking actions where they can have beneficial impact, such as safety and housing.[12]

NIST also states that AI systems can "lead to new services, support, and efficiencies for people, organizations, markets, and society."[13] In addition, the Draft's broad definition of risk—which accounts for potential impacts that are positive,[14] negative, or both—is a helpful way to ensure that the benefits of AI technology get factored into any AI risk assessment.[15] This comprehensive understanding of risk aligns with other standards, as well as other NIST guidance, such as the NISTIR 8286 series.[16]

### B. NIST Should Further Develop Its Discussion of the Complex and Multifaceted Benefits of AI.

While the discussion of AI benefits in the Draft is a good start, NIST should bolster it. As NIST acknowledges, the collective understanding of the impacts of AI systems is still developing.[17] NIST explains that that while "AI benefits and some AI risks are well-known, the AI community is only beginning to understand and classify incidents and scenarios that result in

[12] Draft at 1.

[13] *Id.* at 5.

[14] *Id.* NIST is right to recognize that ISO Guide 73:2009 and IEC/ISO 31010 note that risk can be positive. Still, NIST should recognize that the term "positive risk" may be confusing, since some standards view "risk" as only having adverse impacts.

[15] NIST states that "[T]his framework intends to offer approaches to minimize anticipated negative impacts of AI systems *and* identify opportunities to maximize positive impacts." *Id.*

[16] NIST, NISTIR 8286: Integrating Cybersecurity and Enterprise Risk Management (ERM), at 12 (Oct. 2020), https://doi.org/10.6028/NIST.IR.8286 ("An enterprise that seeks to avoid all cybersecurity risk might stifle innovation or efficiencies to the point where little value would be produced."); NIST, NISTIR 8286A: Identifying and Estimating Cybersecurity Risk for Enterprise Risk Management, at 8 (Nov. 2021), https://doi.org/10.6028/NIST.IR.8286A ("These discussions may also help identify positive risks in the form of opportunities. . . . Understanding that private sector organizations pursue risk as part of their growth strategies and competitive advantage, this aspect should not be forgotten."); NIST, NISTIR 8286B: Prioritizing Cybersecurity Risk for Enterprise Risk Management, at 11 (Feb. 2022), https://doi.org/10.6028/NIST.IR.8286B (discussing considerations of positive risks as enterprise risk management inputs); NIST, Draft NISTIR 8286C: Staging Cybersecurity Risks for Enterprise Risk Management and Governance Oversight, at 19 (Jan. 2022), https://doi.org/10.6028/NIST.IR.8286C-draft ("Each risk aggregation, normalization, and integration activity should identify the impacts of beneficial uncertainty that will accentuate the likelihood of achieving enterprise objectives.").

[17] Draft at 5 (noting that "impacts are not easily foreseeable, and applications are evolving rapidly").

harm."[18] But just as the AI community is only beginning to understand and classify potential harms, the same is true for understanding and classifying benefits, which NIST should devote more resources and attention towards.

The benefits of AI are complex and multifaceted. Like risks, there are different *types* of AI benefits—including benefits to individuals, benefits to enterprises, and benefits to society. For example, AI provides consumers with convenience; algorithmic decision-making helps tailor content to individual users; recommendation algorithms provide consumers with new music or movies based on their viewing or listening habits; and AI can help improve customer service and technical support for consumers. AI can also increase efficiency for enterprises and industry by automating rote tasks to free up human capabilities for more complicated problem solving. Importantly, AI can also help advance broader societal goals such as improving the way people receive services, the efficiency and safety of city operations, and the accessibility of the world to people with disabilities. As mentioned above, AI can deliver cybersecurity and fraud prevention functions, including improving cybersecurity to protect the integrity of critical infrastructure and other networks.[19] Other examples of the tangible benefits that AI can deliver to society

[18] *Id.*

[19] *See, e.g.*, DHS Science and Technology Directorate, *Feature Article: A Secure Environment to Create the Future of Cybersecurity Solutions* (Nov. 16, 2021), https://www.dhs.gov/science-and-technology/news/2021/11/16/feature-article-secure-environment-create-future-cybersecurity-solutions. ("For CISA, the lab will leverage artificial intelligence (AI), machine learning (ML) models, and other advanced data analytics to provide greater situational awareness to inform decision-making regarding the nation's cybersecurity threats."); Business Wire, *C3.ai Digital Transformation Institute Announces Research Awards for AI to Transform Cybersecurity and Secure Critical Infrastructure* (Mar. 24, 2022), https://www.businesswire.com/news/home/20220324005202/en/C3.ai-Digital-Transformation-Institute-Announces-Research-Awards-for-AI-to-Transform-Cybersecurity-and-Secure-Critical-Infrastructure.

abound—AI is helping to increase sustainability;[20] furthering research to detect cancer and develop pharmaceuticals;[21] and helping predict aftershock locations after earthquakes.[22]

There is a growing body of work and empirical research around the positive impacts of AI. For example, in the financial sector, research is looking at AI's potential impacts on access to credit and other financial services.[23] Another important area of research is exploring how AI is used to address the negative outcomes that are often considered to be risks associated with AI—such as bias. As CTIA has explained, while AI could reinforce biases if systems are poorly designed, automated decision-making and algorithms can also help mitigate potential biases and shortcomings in *human* decision-making,[24] which has the potential to yield improved outcomes for consumers, businesses, and society.[25] NIST's involvement in a symposium to address how

[20] *See* AMP Robotics, *AMP Robotics Installs its First Recycling Robots in the United Kingdom and Ireland with Recyco* (Sept. 22, 2021), https://www.amprobotics.com/newsroom/amp-robotics-installs-its-first-recycling-robots-in-the-united-kingdom-and-ireland-with-recyco; Adam Zewe, *Preventing poaching, AI software that predicts poaching hotspots now being deployed to wildlife parks*, Harvard School of Engineering and Applied Sciences (June 16, 2020), https://www.seas.harvard.edu/news/2020/06/preventing-poaching.

[21] *See* Erik Verburg et al., *Deep Learning for Automated Triaging of 4581 Breast MRI Examinations from the DENSE Trial*, 302 Radiology (Oct. 5, 2021), https://pubs.rsna.org/doi/10.1148/radiol.2021203960; Diego Ardila, et al., *End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography*, 25 Nature Medicine 954 (May 20, 2019), https://doi.org/10.1038/s41591-019-0447-x; Dustyn A. Barnette et al., *Lamisil (terbinafine) toxicity: Determining pathways to bioactivation through computational and experimental approaches*, 156 Biochemical Pharmacology 10 (Oct. 2018), https://doi.org/10.1016/j.bcp.2018.07.043.

[22] *See* Phoebe M. R. DeVries et al., *Deep learning of aftershock patterns following large earthquakes*, 560 Nature 632 (Aug. 29, 2018), https://www.nature.com/articles/s41586-018-0438-y.

[23] FinRegLab, *AI in Financial Services*, https://finreglab.org/ai-machine-learning (last visited Apr. 26, 2022).

[24] *See* CTIA AI RMF Concept Paper Comments at 9; Jon Kleinberg et al., *Discrimination in the Age of Algorithms*, 10 J. of Legal Analysis 113, 120 (Apr. 22, 2019), https://academic.oup.com/jla/article/doi/10.1093/jla/laz001/5476086 ("Our central claim, stated in simple form, is that safeguards against the biases of the people who build algorithms, rather than against algorithms per se, could play a key role in ensuring that algorithms are not being built in a way that discriminates (recognizing the complexity and contested character of that term). If we do that, then algorithms go beyond merely being a threat to be regulated; they can also be a positive force for social justice.").

[25] *See, e.g.*, Cass R. Sunstein, *Algorithms, Correcting Biases*, 86 Social Research: An International Quarterly 499, 500 (2019), http://eliassi.org/sunstein_2019_algs_correcting_biases.pdf ("Kleinberg and his colleagues built an algorithm that uses, as inputs, the same data available to judges at the time of a bail hearing, such as prior criminal history and current offense. Their central finding is that *along every dimension that matters, the algorithm does much better than real-world judges*.") (emphasis in original); Jon Kleinberg et al., *Human Decisions and Machine Predictions*, 133 The Quarterly J. of Econ. 237, 241 (Aug. 26, 2017), https://sendhil.org/wp-

technologies such as AI and machine learning "relate to ensuring inclusive economic growth, supporting diversity and financial inclusion, and mitigating risks such as bias and unfairness"[26] is an important step to better understanding these issues and highlighting the many potential benefits of AI.

Accordingly, in the AI RMF, NIST should bolster its discussion of the benefits of AI and incorporate the growing body of work on these benefits. Further, to build on current and existing research, NIST should engage in or encourage additional empirical research around the positive impacts of AI.

## IV. THE AI RMF SHOULD BE A PROCESS-ORIENTED AND POLICY-NEUTRAL RISK MANAGEMENT TOOL.

### A. The AI RMF Should Remain Process-Oriented and Policy-Neutral, Like NIST's Other Successful Risk Management Frameworks.

Drawing on lessons learned from the CSF, NIST is right to build the AI RMF as a process-oriented risk management tool.[27] This will allow the AI RMF to be widely adopted and used across various sectors and use cases, regardless of inevitable changes in technology and

content/uploads/2019/08/Publication-5.pdf ("The algorithm could in principle reduce crime but aggravate racial disparities. Yet the opposite appears to be true in our data: a properly built algorithm can reduce crime and jail populations while simultaneously reducing racial disparities."); Kimberly A. Houser, *Can AI Solve the Diversity Problem in the Tech Industry: Mitigating Noise and Bias in Employment Decision-Making*, 22 Stan. Tech. L. Rev. 290, 352 (2019), https://www-cdn.law.stanford.edu/wp-content/uploads/2019/08/Houser_20190830_test.pdf ("The use of AI in talent-management decisions has shown success in not only creating more successful hires, but in also creating a more diverse slate of candidates and employees. While some companies have embraced these new technologies, others fear that AI may actually cause discriminatory outcomes. As discussed, the phenomena of 'garbage in, garbage out' is real, but can be addressed paying attention to the data sets by using known sources and making sure the sets are balanced and representative of all groups.").

[26] *See, e.g.*, FinRegLab, *Virtual Conferences: Machine Learning Explained and AI in the Economy* (scheduled for April 27-28, 2022), https://finreglab.org/artificial-intelligence-and-the-economy-charting-a-path-for-responsible-and-inclusive-ai. ("The event is designed to address how these technologies relate to ensuring inclusive economic growth, supporting diversity and financial inclusion, and mitigating risks such as bias and unfairness. It will feature presenters and panelists on the cutting edge of researching fairness and explainability in AI, as well as those working to develop policies and frameworks to evaluate and assess the goals of improving the trustworthiness, inclusiveness, and equity of AI deployment.").

[27] Draft at 3 (The AI RMF strives to "[b]e outcome-focused and non-prescriptive. The Framework should provide a catalog of outcomes and approaches rather than prescribe one-size-fits-all requirements.").

law.  Indeed, the CSF's utility and longevity are grounded in its policy neutrality and focus on process rather than specific outcomes.

The AI RMF should similarly not focus on creating specific guidance for certain AI risks. Innovation moves at a rapid pace, and the use cases for AI are constantly expanding.  If the AI RMF becomes a tool intended to tackle specific risks, NIST will constantly be playing a game of catch-up with AI's dynamic and innovative market.  By instead focusing on organizational processes for managing risk, the AI RMF will empower organizations to account for new risks without burdening users with constantly attempting to update responses to narrower use cases.

### B. Recognizing that Trustworthy AI Characteristics and Principles Are Policy-Driven and Nuanced, the AI RMF Should Establish a Flexible Approach that Can Be Adapted to Different Characteristics as Appropriate.

In the Draft, NIST details various technical characteristics, socio-technical characteristics, and guiding principles for AI trustworthiness.[28]  NIST is right to explain that various "characteristics and guiding principles of AI trustworthiness are essential considerations for each function;"[29] however, NIST should reconsider attempting to include definitive classifications of—or policy judgments regarding the implementation of—trustworthy AI characteristics and principles in the AI RMF, as doing so could be counterproductive, is premature, and risks making the AI RMF less usable and more confusing in the long-term.

Application of the trustworthy AI characteristics and principles is inherently a matter of policy and value judgments, and many of the characteristics and principles may be given

[28] The technical characteristics refer to factors that are under the direct control of AI system designers and developers.  These characteristics include accuracy, reliability, robustness, and resilience.  The socio-technical characteristics refer to how AI systems are perceived by individuals, groups, and society at large.  These characteristics include explainability, interpretability, privacy, safety, and how bias is managed.  The guiding principles refer to prioritized societal norms and values. These values include fairness, accountability, and transparency.  *Id.* at 10, 12, 17.

[29] *Id*. at 14.

different weights in different factual contexts.[30] As NIST has explained, the "definitions [of properties like resiliency, reliability, bias, and accountability] vary by author, and they focus on the norms that society expects AI systems to follow."[31] For example, the Draft's Figure 4 maps the AI RMF's trustworthiness characteristics on to other AI policy documents, such as the OECD AI Recommendation, the draft EU AI Act, and Executive Order 13960.[32] Each framework differs in its definition and prioritization of these characteristics, and indeed the draft EU AI Act is currently subject to significant debate. There are twelve trustworthiness characteristics and principles in the AI RMF, which is two more than the EU AI Act and EO 13960, as well as five more than the OECD. The differences between how these characteristics and principles are categorized and prioritized demonstrates the value-laden nature of defining AI trustworthiness. As a practical matter, policy views will differ on the appropriate level of—for example—explainability in AI outcomes or transparency into AI systems, and the approach can also greatly differ by context. The AI RMF will be undercut if NIST attempts to overstep to address these kinds of policy determinations.

Accordingly, in the AI RMF, NIST should recognize that policy judgments underlie many aspects of trustworthy AI and that application of the trustworthy AI characteristics and principles will vary, so it should focus to the greatest extent possible on establishing a flexible, process-based approach to them. Doing otherwise would jeopardize the important policy-neutral qualities that risk management frameworks should embody and could undermine the usefulness of the AI RMF to designers, developers, and users of AI. For example, focusing on detailed

[30] *See e.g.*, *Id.* at 12 ("Guiding principles in the AI RMF taxonomy refer to broader societal norms and values that indicate societal priorities.").

[31] NIST, NISTIR 8312: Four Principles of Explainable Artificial Intelligence, at 1 (Sept. 2021), https://doi.org/10.6028/NIST.IR.8312 ("NISTIR 8312").

[32] Draft at 9.

descriptions of the characteristics and principles may quickly make the AI RMF obsolete in some areas. Further, linking the AI RMF closely to policy judgments that may change over time could threaten the tool's relevance and longevity, making the AI RMF unworkable as divergent regulatory frameworks develop globally. Especially in a dynamic marketplace, the technical capabilities and use cases of AI could change, and policy concerns may grow or lessen in importance.

Additionally, study of trustworthy AI characteristics and principles is still in its early stages, so it would be premature for the AI RMF to include any definitive assessment of them. Indeed, NIST appropriately has separate workstreams in which it is analyzing specific characteristics and principles in more detail—including explainability[33] and bias.[34] As NIST has described with respect to explainability, "[t]he field of explainable AI is an area of active research. Our understanding of these systems and their use will vary as the field grows with new knowledge and data."[35] Similarly, NIST's Special Publication on identifying and managing bias in AI states that it is intended "to surface the salient issues in the challenging area of AI bias, and to provide a *first step* on the roadmap for developing detailed socio-technical guidance for identifying and managing AI bias."[36] In short, because the AI trustworthiness characteristics and principles are still in early phases of development, inclusion of definitive versions of them runs the risk of making the AI RMF obsolete as the concepts inevitably develop and evolve.

[33] *See* NIST, AI Fundamental Research – Explainability (updated Apr. 14, 2022), https://www.nist.gov/artificial-intelligence/ai-fundamental-research-explainability.

[34] *See* NIST, AI Fundamental Research - Free of Bias (updated Apr. 5, 2022), https://www.nist.gov/artificial-intelligence/ai-fundamental-research-free-bias.

[35] NISTIR 8312 at 5.

[36] NIST, NIST Special Publication 1270: Towards a Standard for Identifying and Managing Bias in Artificial Intelligence, at ii (Mar. 2022), https://doi.org/10.6028/NIST.SP.1270 (emphasis added).

At the same time, even if the AI RMF were the appropriate vehicle to detail these trustworthy AI characteristics and principles, the treatment of them in the Draft is oversimplified. Each of the characteristics and principles that NIST identifies is highly complex, but the AI RMF cannot address their complexities within the confines of a risk management tool. For example, NIST's recent Special Publication on identifying and managing bias in AI devotes nine pages to describing the context and terminology of bias in AI,[37] whereas the Draft understandably gives the equivalent section just one paragraph. As a practical matter, the AI RMF cannot attempt to give an issue like bias a full discussion without oversimplification, which could lead to confusion without further explanation.[38] Accordingly, trying to include detailed discussions of these characteristics and principles will ultimately take away from the usability of the AI RMF, and the better alternative is a process-based framework that leaves the more in-depth analyses to separate workflows.

Additionally, each trustworthy AI characteristic and principle is context-specific, and the various characteristics and principles overlap and interrelate. For example, increasing an AI system's explainability or transparency can introduce negative and positive risks. On the one hand, increased explainability or transparency could expose intellectual property and security risks, even if, on the other hand, certain kinds of explainability and transparency could help decrease bias.[39] Additionally, different AI systems, deployed in different contexts, will necessarily need to account for different characteristics and principles. For example, as NIST has recognized, explainability is highly context-dependent because "[t]he audience will strongly

[37] *Id*. at 3-12.

[38] Draft at 12.

[39] NISTIR 8312 at 9-10.

influence the purpose of the explanation and the information it provides. This information will vary according to different groups of people and their role in the system."[40]

These complexities are best addressed in separate and distinct NIST work on these issues, not in the AI RMF, which should provide organizations with a framework to conduct risk management regarding various characteristics and principles, as well as explain that application of these nuanced characteristics and principles requires a flexible approach that may also be driven by separate policy decisions. As discussed above, NIST has already started important work studying many of these characteristics and principles and contributing towards a growing body of work that AI developers and users, as well as other stakeholders, can draw from, including work on explainability and bias. This work should remain separate from the AI RMF, which inevitably would end up oversimplifying analysis of specific characteristics, without benefiting organizations that design, develop, and deploy AI.

Understanding that the congressional directive to create the AI RMF specifically states that the AI RMF should promote trustworthy AI and "establish common definitions and characterizations for aspects of trustworthiness,"[41] CTIA encourages NIST to do so by creating a focused and policy-neutral risk management tool that incorporates trustworthy AI characteristics and principles in a flexible way, recognizing that the specific characteristics and principles that are appropriate for any given AI system are dependent on context and that their implementation ultimately rests on policy decisions.

[40] *Id.* at 6.

[41] 2021 NDAA at § 5301(c) (stating that "[t]he framework shall (1) identify and provide standards, guidelines, best practices, methodologies, procedures and processes for (A) developing trustworthy artificial intelligence systems; (B) assessing the trustworthiness of artificial intelligence systems; and (C) mitigating risks from artificial intelligence systems; (2) establish common definitions and characterizations for aspects of trustworthiness, including explainability, transparency, safety, privacy, security, robustness, fairness, bias, ethics, validation, verification, interpretability, and other properties related to artificial intelligence systems that are common across all sectors").

## V. NIST SHOULD TAKE STEPS TO MAXIMIZE THE PRACTICAL UTILITY AND USABILITY OF THE AI RMF.

### A. NIST Should Hone the Audience for the AI RMF.

In the Draft, NIST defines a broad and diverse range of "key stakeholder[s]" for the AI RMF over four categories: AI system stakeholders, operators and evaluators, external stakeholders, and the general public. The Draft notes that the AI system stakeholders (e.g., business teams, design and development teams; funders; acquisition and procurement teams; internal oversight teams; risk management teams; and compliance teams) are the "primary adopters" of the AI RMF,[42] and it implies that operators and evaluators (e.g., academic, public, and private sector researchers; professional evaluators and auditors; system operators; and expert end users) are also primary adopters of the AI RMF.[43] Beyond these groups, NIST explains that "[i]deally, members of all stakeholder groups would be involved or represented in the risk management process, including those individuals and community representatives that may be affected by the use of AI technologies."[44]

While it is important for NIST to engage with this broad range of stakeholders during the development of the AI RMF, and it is important for designers, developers, and deployers of AI to consider these stakeholders when using AI systems—such a broad set of stakeholders should not all be considered to be the *audience or users* for the AI RMF. For example, NIST should not aim to draft the AI RMF with the average consumer in mind as a *user* of the document, given that AI risk management is highly technical in nature. Likewise, it is not feasible for

[42] Draft at 4.

[43] The Draft explains that "external stakeholders," the third category of stakeholders, are "external to the primary adopters of the AI RMF." *Id.* This implies that the parties listed before the external stakeholders are the primary adopters. This includes operators and evaluators, even though the Draft does not explicitly state that operators and evaluators are primary adopters.

[44] *Id.*

stakeholders like the general public to be incorporated into the process of risk management as it is applied in practice. Rather, for NIST's risk management tool to be most useful, its audience of users should be focused on people and enterprises that are responsible for designing, developing, and/or deploying AI systems.

One justification that has been discussed for the AI RMF having a broader audience is that AI is unique because its risks span beyond an enterprise. While it is true that designers, developers, and users of AI should consider the broader positive and negative impacts on society, the same can be said of other technology and areas of focus. These considerations can and should be part of an enterprise's risk management approach, but the risk management tool to assess these considerations should still be focused on enterprises in order to be practical and useful.

Ultimately, if NIST defines its audience too broadly, the document will not be able to achieve what is necessary for a risk management tool, especially one intended for enterprise use. Focusing on the appropriate audience—while still noting the benefits that the document may have for other, broader stakeholders who are not the intended users—will enable NIST to develop a framework that has the most practical utility, and that can be used flexibly and effectively even if policy goals and requirements for use of AI change over time.

Accordingly, NIST should hone its audience to include only people and enterprises that are responsible for designing, developing, and/or deploying AI systems to foster more effective guidance. NIST can still include a broader set of stakeholders in the development process (e.g., seeking broad feedback and considering the positive and negative impacts associated with AI), without including these broader groups, such as the general public, in the AI RMF's audience of users.

### B. NIST Should Map the AI RMF to Existing Informative References from the Private Sector and Leverage OLIR.

Without endorsing any one approach, NIST should include informative references in the AI RMF. Doing so will allow NIST to align the AI RMF with existing standards work, including but not limited to: the IEEE P7000 suite of standards;[45] IEEE 1012-2016, IEEE Standard for System, Software, and Hardware Verification and Validation;[46] IEEE P3119, Standard for the Procurement of Artificial Intelligence and Automated Decision Systems;[47] and publications from the ISO/IETC committee on AI, SC 42, such as ISO/IEC TF 24027:2021,[48] ISO/IEC CD 42001.2,[49] and ISO/IEC DIS 23894.[50]

To ensure that its informative references are up to date, NIST should also leverage the National Online Informative References ("OLIR") Program, which would help NIST "[t]ake advantage of and foster greater awareness of existing standards, guidelines, best practices, methodologies, and tools for managing AI risks – as well as illustrate the need for additional, improved resources."[51] Utilizing OLIR would help the AI RMF stay up to date by harmonizing

[45] IEEE, *Explore our approved IEEE 7000™ Standards & Projects*, https://ethicsinaction.ieee.org/p7000/ (last visited Apr. 26, 2022).

[46] IEEE, *1012-2016 IEEE Standard for System, Software, and Hardware Verification and Validation* (Sept. 29, 2017), https://ieeexplore.ieee.org/document/8055462.

[47] IEEE, *P3119: Standard for the Procurement of Artificial Intelligence and Automated Decision Systems* (Sept. 23, 2021), https://standards.ieee.org/ieee/3119/10729/.

[48] ISO, *ISO/IEC TF 24027:2021: Information technology — Artificial intelligence (AI) — Bias in AI systems and AI aided decision making* (Nov. 2021), https://www.iso.org/standard/77607.html.

[49] ISO, *ISO/IEC CD 42001.2: Information Technology — Artificial intelligence — Management system,* https://www.iso.org/cms/%20render/live/en/sites/isoorg/contents/data/standard/08/12/81230.html?browse=tc (last visited Apr. 26, 2022).

[50] ISO, *ISO/IEC DIS 23894: Information technology — Artificial intelligence — Risk management,* https://www.iso.org/standard/77304.html (last visited Apr. 26, 2022).

[51] Draft at 3.

with ongoing subject matter expert work on AI standards. Additionally, keeping OLIR up to date would facilitate further and ongoing private sector engagement in the AI RMF.

### C. NIST Should Not Include Prescriptive Guidance or Policy Judgments.

One key to the success of NIST's risk management tools like the CSF is that these tools are not prescriptive. NIST has a long history of defining the "what"—or desired outcomes—but not the "how"—or the specific means for achieving those outcomes. For example, NIST's IoT Device Cybersecurity Capability Core Baseline is an exemplary model for providing this type of guidance, explaining that "[t]he core baseline does not specify how the device cybersecurity capabilities are to be achieved, so organizations who choose to adopt the core baseline for any of the IoT devices they produce, integrate, or acquire have considerable flexibility in implementing it to effectively address needs."[52] NIST should not stray from this approach with the AI RMF.

Accordingly, the next draft of the AI RMF should rephrase any commentary or language that could be construed as being prescriptive guidance. For example, there is a subcategory under "Map" that states: "Benefits of the AI system outweigh the risks, and risks can be assessed and managed. Ideally, this evaluation should be conducted by an independent third party or by experts who did not serve as front-line developers for the system, and who consults experts, stakeholders, and impacted communities."[53] Here, NIST describes not only the "what"—that organizations evaluate risks versus benefits—but also the "how"—recommending that an independent third party conduct this assessment. While some AI systems may call for an independent third-party assessment, some may not. Organizations should be allowed to make

[52] *See, e.g.*, NIST, NISTIR 8259A: IoT Device Cybersecurity Capability Core Baseline at 3 (May 2020), https://doi.org/10.6028/NIST.IR.8259A.

[53] Draft at 16.

the decision of how best to perform the assessment, based on various factors, and NIST's AI RMF should not take a position on the best means by which to achieve the assessment.

Similarly, and as discussed above, NIST should refrain from making policy judgments in the AI RMF. As currently drafted, NIST states that "[i]ndividual human operators and their organizations should be answerable and held accountable for the outcomes of AI systems, particularly adverse impacts stemming from risks."[54] This type of statement has the potential to stray into a policy judgment on who should bear liability rather than staying focused on the general principle of accountability. NIST should remove this and other similar statements.

In general, NIST should ensure that the AI RMF does not layer in prescriptive guidance or value judgments. NIST guidance is successful because it does not prescribe the means of accomplishing goals and does not create policy, so in the AI RMF, NIST should remove any existing language that would stray from this approach.

### D. NIST Should Discuss How Application of the Framework Will Vary Between Low-Risk and Higher-Risk AI Systems.

Different kinds of AI technology are being utilized by and integrated with a wide array of different use cases. While CTIA recognizes that some use cases may have significant impacts, many other use cases are relatively low risk and benign. For low-risk cases, NIST should recognize that the application of the AI RMF will not be complicated, and will vary from higher risk use cases.[55]

[54] *Id.* at 12-13 (NIST repeats this statement twice).

[55] For example, as was explained at the recent AI RMF Workshop, "AI is being integrated into all sorts of products and services now that . . . people would [not] reasonably expect would require being run through the full Framework. So, for instance, [some] word processing software [uses AI] to recognize what . . . font was used." *Building the NIST AI Risk Management Framework: Workshop #2*, at 46:40 – 47:02 (Mar. 29, 2022), https://www.nist.gov/news-events/events/2022/03/building-nist-ai-risk-management-framework-workshop-2 (testimony of Christian Troncoso, Senior Director, BSA: The Software Alliance).

Further, with the forthcoming Practice Guide, NIST should help standardize the risk management expectations around commonplace, low-risk AI use cases. NIST notes in the Draft that a Practice Guide will be released for comment, which will include examples and practices for assistance in adopting the AI RMF.[56] CTIA looks forward to collaborating with NIST on this document. Along with other resources in the AI Resource Center, NIST could leverage the Practice Guide to highlight the flexible nature of the AI RMF and show how the AI RMF may be applied to low-risk AI systems, versus higher-risk applications.

## VI. THE AI RMF WILL SERVE AS AN IMPORTANT FOUNDATION FOR FUTURE EFFORTS TO IMPROVE AI RISK MANAGEMENT, BUT NIST SHOULD TAKE THESE EFFORTS ONE STEP AT A TIME.

### A. NIST Should Finalize the AI RMF Before Beginning Work on Profiles.

The Draft indicates that NIST plans to develop AI RMF Profiles as part of the overall development process of the AI RMF.[57] CTIA agrees that profiles are important tools to accompany risk management frameworks. The approach of customizing a framework for specific risk or use case profiles has been successfully used in deploying NIST's CSF.[58] Under this approach, individual organizations or sectors are empowered to create profiles for particular implementation scenarios based on their needs and unique risks. For example, in the context of the CSF, the Communications Sector has developed CSF implementation guidance for each of

[56] Draft at i ("That Practice Guide which will be released for comment includes additional examples and practices that can assist in using the AI RMF. The Guide will be part of a NIST AI Resource Center that is being established.").

[57] "Development of profiles is deferred until later drafts of the AI RMF are developed with the community. NIST welcomes contributions of AI RMF profiles. These profiles will inform NIST and the broader community about the usefulness of the AI RMF and likely lead to improvements which can be incorporated into future versions of the framework." Draft at 20.

[58] *See* NIST, *Cybersecurity Framework: Questions and Answers* (updated Feb. 24, 2022), https://www.nist.gov/cyberframework/frequently-asked-questions/framework-basics#basics (explaining that "[t]he Framework is guidance," and "[i]t should be customized by different sectors and individual organizations to best suit their risks, situations, and needs").

the five key segments of the industry, including wireless.[59] In addition, CSF profiles have been created to address specific technologies (e.g., NIST's PNT profile)[60] and threats (e.g., NIST's Ransomware Profile).[61]

With that said, it is important that NIST finalize the AI RMF *before* work on any profiles is started. Specifically, before NIST or others create profiles to customize use of the AI RMF, NIST should finalize the AI RMF—after appropriate comment periods—in order to better collaborate with the broader AI community and build profiles that are efficient and effective.

### B. NIST Should Not Attempt to Evaluate the Effectiveness of the AI RMF Until It Is Adopted and Used in the Field.

Similarly, the Draft indicates that there will be a section of the AI RMF dedicated to "Effectiveness of the AI RMF," but that development of this section will be deferred to later drafts.[62] While CTIA supports encouraging organizations to conduct internal self-evaluations of their risk management practices, NIST should be careful not to attempt to develop procedures for evaluating the effectiveness of a document that has not yet been finalized. For the initial iteration of the AI RMF, it may be most appropriate to capture the internal self-assessment work in the "Govern" Function.

[59] *See generally* CSRIC IV, Working Group 4, Cybersecurity Risk Management and Best Practices, Final Report, at 4-5 (March 2015), https://transition.fcc.gov/pshs/advisory/csric4/CSRIC_IV_WG4_Final_Report_031815.pdf.

[60] NIST, *Responsible Use of Positioning, Navigation and Timing Services*, https://www.nist.gov/pnt (last visited Apr. 26, 2022).

[61] NIST, NISTIR 8374: Ransomware Risk Management: A Cybersecurity Framework Profile (Feb. 2022), https://doi.org/10.6028/NIST.IR.8374.

[62] Draft at 20.

## VII. CONCLUSION.

CTIA is pleased to help NIST as it develops the AI RMF and to continue to collaborate to make the document practical and robust to support this important area of emerging technology.

Respectfully submitted,

*/s/ Melanie K. Tiano*_
Melanie K. Tiano
Assistant Vice President, Cybersecurity and Privacy

Thomas K. Sawanobori
Senior Vice President and Chief Technology Officer

John A. Marinho
Vice President, Technology and Cybersecurity

Avonne S. Bell
Director, Connected Life

**CTIA**
1400 16th Street, NW, Suite 600
Washington, DC 20036
202-736-3200
www.ctia.org

April 29, 2022