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

In the Matter of

AI Risk Management Framework	)	Second Draft
	)	
NIST AI Risk Management Framework Playbook	))	Draft

COMMENTS OF CTIA

Thomas K. Sawanobori
Senior Vice President and Chief Technology
Officer

John A. Marinho
Vice President, Technology and Cybersecurity

Avonne S. Bell
Director, Connected Life

Justin C. Perkins
Manager, Cybersecurity and Policy

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

September 29, 2022

Table of Contents

I. INTRODUCTION AND SUMMARY.

CTIA[1] welcomes the opportunity to continue to engage with NIST as it develops the AI Risk Management Framework ("AI RMF")[2] and is pleased to submit comments on the AI Risk Management Framework: Second Draft ("Second Draft") and the draft NIST AI Risk Management Framework Playbook ("Draft Playbook"). The Second Draft builds on the foundation of NIST's prior guidance documents, and it makes several important improvements from the initial public draft ("Initial Draft"). With these comments, CTIA highlights those improvements and recommends that NIST make further changes to the Second Draft and Draft Playbook to ensure that these documents are policy-neutral, flexible, and risk-based tools that can be used by a wide range of stakeholders. Specifically:

- In the AI RMF, NIST should (1) expand on the benefits of AI; (2) continue to ensure that any discussion of trustworthiness characteristics is generalized and flexible—to help the AI RMF remain policy-neutral; (3) remove any remaining prescriptive language; (4) further refine the audience of the AI RMF; and (5) directly incorporate existing risk management standards and tools.
- In the Playbook, NIST should (1) ensure that it is flexible, risk-based, policy-neutral, and not unduly prescriptive; (2) clarify that the Playbook is meant to provide potential tools for organizations that design, develop, and deploy AI; and (3) adjust the Draft Playbook to account for a wide range of organizations that will utilize AI, especially for those using low-risk AI systems.

[1] CTIA® (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] During the development process, CTIA has submitted comments in response to NIST's initial Request for Information and the *AI Risk Management Framework Concept Paper*, as well as the initial public draft ("Initial Draft"); participated in the October 2021 Kickoff Workshop; and facilitated a meeting between NIST's AI RMF development team and CTIA members.

- Regarding the overall AI RMF development process, NIST (1) should not promote the development of AI RMF profiles until the AI RMF is complete; and (2) should align the Playbook's update schedule with that of the AI RMF.

II. NIST'S SECOND DRAFT TAKES IMPORTANT STEPS TO ENSURE THAT THE AI RMF IS A POLICY-NEUTRAL TOOL FOR ORGANIZATIONS TO MANAGE RISKS AND MAXIMIZE BENEFITS ASSOCIATED WITH AI.

A. The Second Draft Emphasizes that the AI RMF Is Voluntary, Risk-Based, and Flexible—Attributes that Will Best Facilitate Innovation and Beneficial Uses of AI.

The Second Draft generally retains the Initial Draft's "Attributes,"[3] which notably state that the AI RMF strives to: "[b]e risk-based, resource-efficient, pro-innovation, and voluntary[]" and "[b]e outcome-focused and non-prescriptive," noting that the AI RMF "should provide a catalog of outcomes and approaches rather than prescribe one-size-fits-all requirements."[4] Importantly, the Second Draft maintains an emphasis on the voluntary and flexible nature of the AI RMF. For example, it states that "[t]he AI RMF is *intended for voluntary use* to address risks in the design, development, use, and evaluation of AI products, services, and systems."[5] The Second Draft is also focused on process and is intended to be applicable to a wide range of AI implementations.[6] NIST explains that the AI RMF is not a "checklist and it is not intended to be used in isolation."[7] Further, NIST grounds the Second Draft in principles of risk management by cautioning users to avoid attempts to eliminate risk entirely. NIST encourages fostering a "risk mitigation culture," as opposed to attempting to eliminate risk entirely, because "incidents and

[3] NIST, AI Risk Management Framework: Second Draft at 4 (Aug. 18, 2022), https://www.nist.gov/system/files/documents/2022/08/18/AI_RMF_2nd_draft.pdf ("Second Draft").
[4] *Id*.
[5] *Id.* at i. (emphasis in original); *see also id.* at 2 ("The AI RMF is a voluntary framework seeking to provide a flexible, structured, and measurable process to address AI risks prospectively and continuously throughout the AI lifecycle.").
[6] *Id*. ("The Framework describes a process for managing AI risks across a wide spectrum of types, applications, and maturity – regardless of sector, size, or level of familiarity with a specific type of technology.").
[7] *Id*.

failures cannot be eliminated."[8] CTIA agrees that risk elimination efforts can be "counterproductive" by fostering unrealistic expectations and putting undue stress on organizational resources.[9]

These attributes are critically important to the success of the AI RMF, as they are the signature attributes of important NIST efforts, including the Cybersecurity Framework ("CSF"). NIST should continue to follow its tried-and-trusted path with the AI RMF. In particular, because AI technology is developing rapidly, with a wide range of varied use cases, a voluntary, flexible, and risk-based tool will create an environment where innovation is encouraged.

B. The Second Draft Expands on the Benefits of AI, Which NIST Should Continue to Explore and Assess.

AI systems and technologies promise enormous benefits, including in the telecommunications industry.[10] For example, the wireless sector can implement AI to engage in real-time network threat detection, combat fraud, improve customer service, and help build 5G networks, among many other activities that provide tremendous benefits.[11]

[8] *Id*. at 10.

[9] *Id*.

[10] The Federal Communications Commission's ("FCC") industry-led Technological Advisory Council ("TAC") has explained that AI could broadly apply across the industry. *See* TAC, 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] *See* Jason Porter & Mazin Gilbert, *From Raw Data to Real Solutions: How We're Applying AI*, AT&T (May 24, 2019), https://about.att.com/innovationblog/2019/05/applied_ai.html; 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 Sept. 20, 2022); AWS, *Using Technology to Improve Personal Connections*, https://aws.amazon.com/machine-learning/customers/innovators/t_mobile/, (last visited Sept. 20, 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. 9, 2021), https://www.qualcomm.com/news/onq/2021/11/10/qualcomm-power-efficient-ai-making-technology-more-sustainable.

While the Initial Draft started the conversation about the positive impacts of AI, CTIA has encouraged NIST to devote more attention to understanding and assessing the full range of AI benefits.[12] The Second Draft takes steps in this direction. For example, the Second Draft rightly explains that use of the AI RMF could allow users to increase and enhance benefits to society through the use of AI.[13] Additionally, the discussion of the Map function counsels organizations to assess both risks and benefits during the risk management process.[14] The Second Draft further explains that "[s]ince AI systems can make sense of information more quickly and consistently than humans, they are often deployed in high-impact settings as a way to make decisions fairer and more impartial than human decision-making, and to do so more efficiently."[15] CTIA supports NIST's inclusion of further discussion about the benefits of AI.

But still, there is more work that NIST can and should do to study and incorporate the wide range of benefits into the AI RMF. *First,* while it is important to understand the potential negative impacts and have tools to identify and mitigate negative risks, NIST should also provide more specific examples and guidance about considering benefits in the risk management analysis. For example, NIST could create categories of benefits—such as benefits to individuals, benefits to organizations, and benefits to society—for organizations to consider in their risk management approach.

[12] *See* Comments of CTIA, NIST AI Risk Management Framework: Initial Draft at 5-8 (Apr. 29, 2022) ("Comments of CTIA").

[13] Second Draft at 3 ("Using the AI RMF may reduce the likelihood and degree of negative impacts and increase the benefits to individuals, groups, communities, organizations, and society."); *id.* at 7 ("While risk management processes address negative impacts, this framework offers approaches to minimize anticipated negative impacts of AI systems and identify opportunities to maximize positive impacts."); *id.* at 11 ("Increasing the breadth and diversity of stakeholder input throughout the AI lifecycle can enhance opportunities for identifying AI system benefits and positive impacts, and increase the likelihood that risks arising in social contexts are managed appropriately.").

[14] *Id*. at 20.

[15] *Id*. at 12.

Second, NIST should draw on existing and developing research on the benefits of AI, including the potential for AI—properly managed—to be used to fight against negative outcomes, such as bias and discrimination. For example, AI has the potential to increase access to financial services, and FinRegLab, an organization with which NIST has partnered on various AI efforts, is researching the responsible and fair use of AI and machine learning ("ML") in financial services.[16] Additionally, AI can aid in sustainability efforts[17] and assist medical researchers.[18] Also, with appropriate controls, AI can help reduce bias in human decision-making.[19] NIST should draw on this research in the AI RMF.

[16] FinRegLab, *AI in Financial Services*, https://finreglab.org/ai-machine-learning.

[17] *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 John. A. Paulson Sch. of Eng'g and Applied Scis. (June 16, 2020), https://www.seas.harvard.edu/news/2020/06/preventing-poaching.

[18] *See* Erik Verburg et al., *Deep Learning for Automated Triaging of 4581 Breast MRI Examinations from the DENSE Trial*, 302 Radiology 29 (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 Med. 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.

[19] *See Mitigation of AI/ML Bias in Context*, NCCoE, https://www.nccoe.nist.gov/projects/mitigating-aiml-bias-context ("Automated decision-making is appealing because artificial intelligence (AI)/machine learning (ML) systems produce more consistent, traceable, and repeatable decisions compared to humans . . . ."). *See also, e.g.*, Cass R. Sunstein, *Algorithms, Correcting Biases*, 86 Soc. Rsch.: An Int'l Q. 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 Q. J. of Econ. 237, 241 (Feb. 2018), https://sendhil.org/wp-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."); Ashesh Rambachan & Jonathan Roth, *Bias in, bias out? Evaluating the folk wisdom*, 1st Symposium on Foundations of Responsible Computing (FORC 2020) 6:1 (May 18, 2020), https://scholar.harvard.edu/files/jroth/files/lipics-forc-2020-6.pdf; Kimberly 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.").

Third, NIST should encourage more empirical research into studying the benefits of AI. To this end, NIST should consider facilitating work to explore the positive benefits of AI, similar to its research and work into AI risks.

C. The Second Draft Makes Progress Towards Being Policy-Neutral, but the Discussion of Trustworthiness Characteristics Still Needs Improvements.

Like the CSF, the AI RMF should be process-oriented as well as technology- and policy-neutral. While the Second Draft makes progress towards this goal, especially with respect to the discussion of trustworthiness characteristics, there is more that NIST should do to address these value-laden issues without undermining neutrality.

1. The AI RMF Should Be Policy-Neutral and Process-Focused.

The CSF continues to be widely used and effective because it is centered around process,[20] and the AI RMF should follow this approach. Ensuring that any risk management tool is focused on high-level process and does not attempt to dictate substantive requirements is highly important for ensuring continued relevance and longevity, given the rapid development of both the AI technological and legal landscape. Indeed, innovation in AI technology is moving at a rapid pace, with the use cases for AI constantly expanding.[21] Having a process-focused tool will help NIST's guidance not to lag behind AI's dynamic market. The AI legal and regulatory landscape is also evolving, so it is important to have a policy-neutral tool that organizations can

[20] NIST, Framework for Improving Critical Infrastructure Cybersecurity Version 1.1 at v (Apr. 16, 2018), https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf ("CSF 1.1") (explaining that "the Framework will help an organization to align and prioritize its cybersecurity activities with its business/mission requirements, risk tolerances, and resources.").

[21] *See, e.g.*, Second Draft at i. ("AI research and development, as well as the standards landscape, is evolving rapidly. For that reason, the AI RMF and its companion documents will evolve over time and reflect new knowledge, awareness, and practices."); Joe McKendrick, *AI Adoption Skyrocketed Over the Last 18 Months*, Harvard Business Review (Sept. 27, 2021), https://hbr.org/2021/09/ai-adoption-skyrocketed-over-the-last-18-months; Thomas Helfrich, *Why Robotics And Artificial Intelligence Are The Future Of Mankind*, Forbes (May 31, 2022), https://www.forbes.com/sites/forbestechcouncil/2022/05/31/why-robotics-and-artificial-intelligence-are-the-future-of-mankind/?sh=160036291689.

consult and use, as applicable, across jurisdictions and in the context of the legal framework that may apply to any given AI system. For example, international frameworks for addressing responsible AI are developing, including the EU's proposed AI Act.[22] There are also many federal and state initiatives—some new or proposed, others more mature—that impact AI. Examples include:

- In 2020, Executive Order 13960 established principles for federal agency use of AI.[23]
- Federal laws, such as the Fair Credit Reporting Act and the Equal Credit Opportunity Act,[24] apply to the use of algorithms when used to make certain decisions covered under those laws.
- The Federal Trade Commission ("FTC") has provided guidance on AI and, most recently, the agency issued an Advanced Notice of Proposed Rulemaking that includes questions about possible new policy approaches for "automated decision-making systems," among other privacy and data security-related topics.[25]
- Some states have passed bills that create tasks forces that will study issues relating to the use of AI.[26]
- Other states have established, under omnibus privacy statutes, a consumer's right to opt-out when covered businesses engage in certain automated decisionmaking activities.[27]

Given the developing and dynamic AI legal environment, NIST should continue to keep the AI RMF policy-neutral and avoid prescribing specific approaches to specific risks.

[22] *See, e.g.*, Proposal for a Regulation of the European Parliament and the of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, COM (2021) 206 final (April 4, 2021), https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206.

[23] Exec. Order No. 13,960, 85 Fed. Reg. 78,939 (Dec. 12, 2020), https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government.

[24] *See, e.g.*, Fair Credit Reporting Act, 15 U.S.C. §§ 1681-1681x; Equal Credit Opportunity Act, 15 U.S.C. §§ 1691-1691f.

[25] *Trade Regulation Rule on Commercial Surveillance and Data Security*, Advance Notice of Proposed Rulemaking, 87 Fed. Reg. 51,273, 51,283-84 (Aug. 22, 2022), https://www.federalregister.gov/documents/2022/08/22/2022-17752/trade-regulation-rule-on-commercial-surveillance-and-data-security.

[26] *See, e.g.*, S.B. 78, 2021 Reg. Sess. (Ala. 2021); S.B. 22-113, 73rd Gen. Assemb., 2022 Reg. Sess. (Colo. 2022); H.B. 0645, 102nd Gen. Assemb., 2021 Reg. Sess. (Ill. 2021); H. 410, 2022 Reg. Sess. (Vt. 2022).

[27] *See, e.g.*, Virginia Consumer Data Protection Act ("VCDPA"), Va. Code Ann. § 59.1-575 et seq.; Colorado Privacy Act ("CPA"), Colo. Rev. Stat. § 6-1-1301 et seq.; Connecticut Data Privacy Act ("CTDPA"), S.B. 6, 2022 Reg. Sess. (Conn. 2022).

Generally speaking, and with caveats noted in the next section, CTIA believes the Second Draft helpfully retains much of the Initial Draft's commitment to a process-focused and policy- and technology-neutral approach, which CTIA supports. The Second Draft explains that the AI RMF is meant to "describe[] a process for managing AI risks across a wide spectrum of types, applications, and maturity – regardless of sector, size, or level of familiarity with a specific type of technology."[28] It also rightly clarifies that "[t]he AI RMF is not a compliance mechanism. It is law- and regulation-agnostic, as AI policy discussions are live and evolving. While risk management practices should incorporate and align with applicable laws and regulations, this document is not intended to supersede existing regulations, laws, or other mandates."[29]

2. NIST Should Continue to Refine the Discussion of Trustworthiness Characteristics to Ensure that It Is Flexible, Does Not Categorically Define Characteristics, and Does Not Attempt to Judge Tradeoffs.

CTIA does recommend, however, that NIST revise certain aspects of the discussion of trustworthiness characteristics to ensure that they are not categorically defined and that potential tradeoffs between characteristic are not pre-judged—to avoid policy judgments and maintain flexibility in the implementation and inevitable evolution of the RMF. Because identifying characteristics of trustworthy AI is inherently value-laden, it is critical that any discussion of these characteristics is flexible, general, and policy-neutral. In response to the Initial Draft, CTIA suggested that the proposed discussion of trustworthy AI characteristics and principles was overly complex and definitive, and in some instances included policy judgments regarding certain characteristics or principles played out in AI systems.[30] CTIA believes that the Second Draft makes significant improvements, including adding a discussion of specific trustworthiness

[28] Second Draft at 2.

[29] *Id*.

[30] *See* Comments of CTIA at 9-13.

characteristics that is more streamlined and accessible for users. In addition, the Second Draft helpfully emphasizes that these characteristics are nuanced and will differ from one application to the next,[31] and makes clear that not all trustworthiness characteristics will be relevant to all AI systems and that they may not need to be strictly "measured" in all circumstances, depending on the context.[32]

There are still a few areas in which NIST should continue to refine its discussion of AI trustworthiness characteristics to ensure that the AI RMF remains policy-neutral and flexible. *First,* to account for the fact that organizations may take differing approaches even to very similar AI systems, the trustworthiness characteristics should explicitly highlight how each characteristic is context dependent. The "transparent and accountable" characteristic description does this well, explaining for example that "[t]he relationship between risk and accountability associated with AI and technological systems more broadly differs across cultural, legal, sectoral, and societal contexts."[33] Such language should be expanded to make clear that the characteristic will vary from system to system based on this type of context, and NIST should include similar language for the other trustworthiness characteristics.

Second, NIST should not attempt to categorically identify or define *all* trustworthiness characteristics. As drafted, the Second Draft states definitively that: "Trustworthy AI ***is***: valid

[31] Second Draft at 4 ("The decision to commission or deploy an AI system should be based on a contextual assessment of trustworthiness characteristics and the relative risks, impacts, costs, and benefits, ..."); *id.* at 11 ("Addressing AI trustworthy characteristics individually will not assure AI system trustworthiness, and ***tradeoffs are always involved***. Trustworthiness is greater than the sum of its parts. Ultimately, it is a social concept, and ***the characteristics listed in any single guidance document will be more or less important in any given situation to establish trustworthiness***.") (emphasis added); *id.* ("Human judgment must be employed when deciding on the specific metrics related to AI trustworthy characteristics and the precise threshold values for their related metrics.").

[32] *Id*. at 23 (For example, Measure 1.1 explains: "Approaches and metrics for quantitative or qualitative measurement of the most significant risks, identified by the outcome of the Map function, including context-relevant measures of trustworthiness are identified and selected for implementation. The risks or *trustworthiness characteristics that will not be measured* are properly documented.") (emphasis added).

[33] *Id*. at 15.

and reliable, safe, fair and bias is managed, secure and resilient, accountable and transparent, explainable and interpretable, and privacy-enhanced."[34] While CTIA agrees that a general and streamlined articulation of the common characteristics that are emerging is helpful, NIST should be less absolute about identifying a single set of characteristics or definitions of the characteristics. NIST should consider noting that different frameworks have identified different characteristics, and that while NIST's list seeks to group these characteristics into general categories, new or different categories may exist and/or emerge. Further, NIST should clarify that definitions for these characteristics vary, as NIST has recognized in its more specific work regarding certain trustworthiness characteristics.[35]

Third, the RMF should avoid overly prescriptive guidance for how the characteristics should be implemented in practice, and instead rely on a risk-based approach, while leaving more practical guidance to other workstreams. The AI RMF should not attempt to assess the tradeoffs associated with addressing the various trustworthiness characteristics; that work requires policy judgments and specific risk assessments. To this end, NIST uses the right approach when it cites to either applicable standards or applicable NIST documents (e.g., ISO/IEC standards, NISTIR 8312, NIST SP 1270, NISTIR 8312, NISTIR 8367, and the NIST Privacy Framework) in its discussions of the various characteristics. NIST, however, should remove standalone guidance within the AI RMF itself. The section on the "valid and reliable" characteristic states, for example, that "[a]ccuracy measurements *should always* be paired with clearly defined test sets and details about test methodology; both should be included in

[34] *Id*. at 10 (emphasis added).

[35] Indeed, NIST in its other AI work has stated that "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." NIST, NISTIR 8312, Four Principles of Explainable Artificial Intelligence at 1 (Sept. 2021), https://doi.org/10.6028/NIST.IR.8312 ("NISTIR 8312").

associated documentation."[36] These discussions are better left to other workstreams, including NIST's work on specific trustworthiness characteristics, and the AI RMF should only contain references to that more specific guidance, without including their specific suggestions. Finally, NIST should explicitly note that each of these references includes voluntary and flexible guidance that does not create requirements for users.

D. Removing Prescriptive Language Improves the Second Draft, but NIST Should Remove Any Remaining Prescriptive Language Before the Document Is Finalized.

In the context of risk management guidance, NIST has a long history of defining the "what" (i.e., the desired outcomes), but not the "how" (i.e., the specific way to achieve those outcomes). NIST's IoT Baseline is an exemplary model for providing this type of guidance; it sets outs to create a baseline, but it acknowledges that not every element of the baseline will be applied similarly—if at all—depending on context.[37] CTIA has urged NIST not to stray from that approach with the AI RMF, and for the most part, the Second Draft's AI RMF Core is appropriately outcome-focused and not prescriptive. As one example, in the Initial Draft, CTIA identified a subcategory that could have been read to prescribe the "how" instead of simply identifying the "what,"[38] and NIST rightly removed this subcategory in the Second Draft.

[36] Second Draft at 13 (emphasis added).

[37] NIST, NISTIR 8259A, IoT Device Cybersecurity Capability Core Baseline at 3 (May 2020), https://doi.org/10.6028/NIST.IR.8259A ("The core baseline's role is as a default for minimally securable devices. However, device cybersecurity capabilities will often need to be added or removed from an IoT device's design, integration, or acquisition to best address an organization's common cybersecurity risks.").

[38] The fourth Map category had stated, "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." NIST, AI Risk Management Framework: Initial Draft at 16 (Mar. 17, 2022), https://www.nist.gov/system/files/documents/2022/03/17/AI-RMF-1stdraft.pdf ("Initial Draft").

As NIST finalizes the AI RMF, it should carefully review the AI RMF Core to remove any remaining prescriptive language, including—but not limited to— the items listed in the table below.

AI RMF Core Subcategory	CTIA Recommendation
GOVERN 3.1: "Decision-making related to mapping, measuring, and managing AI risks throughout the lifecycle is informed by a demographically and disciplinarily diverse team ***including internal and external personnel***. Specifically, teams that are directly engaged with identifying design considerations and risks include a diversity of experience, expertise, and backgrounds to ensure AI systems meet requirements beyond a narrow subset of users."[39]	It would be more appropriate for NIST to explain these teams *may include* external personnel.
MEASURE 1.3: "Internal experts who did not serve as frontline developers for the system and/or independent assessors are involved in regular assessments and updates. Domain experts, users, and external stakeholders and affected communities are consulted in support of assessments."[40]	This subcategory presupposes that certain stakeholders cannot perform adequate assessments, or that certain stakeholders are required for adequate assessments. Its conclusions may be accurate in some scenarios, but in others, it would be appropriate for developers to conduct robust assessments. The individual(s) that conduct assessments should be determined based on the specific context of the organization and AI system.
MEASURE 2.8: "AI model is explained, validated, and documented. AI system output is interpreted within its context and to inform responsible use and governance."[41]	Different AI systems will call for different explainability needs. NIST's subcategory should include "as appropriate," to account for this range.

E. The Second Draft's Discussion About How the AI RMF May Be Used in the Context of Low-Risk AI Systems Helps to Illustrate the Varied Use Cases and Risk Management Approaches for AI.

The Second Draft includes discussion of low-risk AI systems, consistent with CTIA's past advocacy.[42] Because of the wide array of AI uses cases—which are expanding daily—it is important for NIST to note that different AI systems will have different risk profiles. While

[39] Second Draft at 20 (emphasis added).
[40] *Id*. at 23.
[41] *Id*. at 24.
[42] Comments of CTIA at 18 ("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.").

some AI systems may be quite impactful, other systems will be low-risk and benign; an organization's risk management approach will vary based on these differing risk profiles. The Second Draft clearly explains that organizations should prioritize their resources for high-risk systems while decreasing prioritization for low-risk systems.[43] It also includes a helpful discussion of an organization's risk tolerance. NIST notes that "[w]hile the AI RMF can be used to prioritize risk, it does not prescribe risk tolerance. . . . Risk tolerance and the level of risk that is acceptable to organizations or society are highly contextual and application and use-case specific."[44] The Draft Playbook also discusses low-risk AI systems, as GOVERN 5.2 and GOVERN 6.1 both recommend that organizations should direct larger allocations of resources to high risk-systems and fewer resources to low-risk systems.[45] NIST is right to explain that low-risk AI systems will require a less rigorous approach, as well as fewer organizational resources for oversight. This distinction is fundamental to risk management and will allow users to conserve AI-related resources and direct them to higher-priority use cases.

III. NIST SHOULD MAKE ADDITIONAL IMPROVEMENTS TO MAKE THE AI RMF AN EFFECTIVE AND PRACTICAL RISK MANAGEMENT TOOL.

A. NIST Should Make Clear that AI RMF Users Should Be Individuals and Organizations that Design, Develop, and Deploy AI.

In prior advocacy, CTIA asked NIST to hone the audience for the AI RMF, which was initially construed as including AI system stakeholders, operators and evaluators, external

[43] Second Draft at 9 ("When applying the AI RMF, risks which the organization determines to be highest for AI systems and contexts call for the most urgent prioritization and most thorough risk management process. Doing so can mitigate risks and harms to enterprises and individuals, communities, and society. In some cases where an AI system presents the highest risk – where negative impacts are imminent, severe harms are actually occurring, or catastrophic risks are present – development and deployment should cease in a safe manner until risks can be sufficiently mitigated. Conversely, the lowest-risk AI systems and contexts suggest lower prioritization.").

[44] *Id*. at 9-10.

[45] *See* NIST, *AI RMF Playbook: GOVERN 5.2*, https://pages.nist.gov/AIRMF/govern-5/2004/05/02/post-feedback.html; NIST, *AI RMF Playbook: GOVERN 6.1*, https://pages.nist.gov/AIRMF/govern-6/2004/06/01/guideline-for-third-party-systems.html.

stakeholders, and the general public.[46] 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 the *audience* or *users* for the AI RMF. Rather, for NIST's risk management tool to be most useful, its audience of users should be focused on people and organizations that are responsible for designing, developing, and/or deploying AI systems. In the Second Draft, NIST has taken steps in the right direction, but before the AI RMF is finalized, further clarification about its intended audience and users will be important. Ultimately, if NIST attempts to define the audience too broadly (e.g., looking beyond people and enterprises that are responsible for designing, developing, and/or deploying AI systems to include the general public and consumers/end users), it will undermine the effectiveness and value of the tool.

Specifically, in the Second Draft, NIST rightly clarifies that: (1) the AI RMF is intended to be used by "AI actors," meaning "those who play an active role in the AI system lifecycle, including organizations and individuals that deploy or operate AI,"[47] and (2) the AI RMF's "primary audience" includes test, evaluation, verification, and validation ("TEVV") AI actors that design, develop, deploy, and acquire AI systems.[48] Additionally, the primary audience includes "those with responsibilities to commission or fund an AI system and those who are part of the enterprise management structure governing the AI system lifecycle."[49]

[46] Comments of CTIA at 14-15.
[47] Second Draft at 4 (quoting OECD Library, Artificial Intelligence in Society (June 11, 2019), https://www.oecd-ilibrary.org/science-and-technology/artificial-intelligence-in-society_eedfee77-en).
[48] *Id*. at 6. NIST notes that these actors "are especially likely to benefit from the Framework." *Id.*
[49] *Id*. at 5-6.

Nevertheless, the scope of the Second Draft is still overly broad due to NIST's "People & Planet" audience dimension, which includes "end users and affected entities."[50] NIST should not aim for the AI RMF's guidance to be applicable to the average consumer, and consumers should not be considered users. While it is true that designers, developers, and users of AI should consider the broader impacts on society—both positive and negative—the same can be said of other technologies and areas of focus. These considerations should be part of an enterprise's risk management process, but the risk management tool for assessing such considerations should still target enterprises in order to be practical and useful. To increase the practicality of the AI RMF, NIST should explain that the AI RMF is not intended to be *used* by end-users or affected entities, and narrow its audience so that the AI RMF achieves its targeted goal: helping *organizations* manage AI risk.

Further, it is not appropriate for the AI RMF to suggest that AI actors must solicit and account for community feedback in every, or even most, AI deployments, regardless of the risk profile—as is presupposed in the Second Draft. For example, the Second Draft states: "The People & Planet dimension of the AI lifecycle . . . presents an additional AI RMF audience: *end-users* or affected entities who play an important consultative role to the primary audiences. Their insights and input equip others to analyze context, identify, monitor and manage risks of the AI system by providing formal or quasi-formal norms or guidance."[51] Two subcategories similarly envision engagement with outside stakeholders: (1) "GOVERN 5.2: Mechanisms are established to enable AI actors to regularly incorporate adjudicated stakeholder feedback into system design

[50] *Id*. at 6-7.

[51] *Id*. at 6 (emphasis added). Later in the Draft, NIST adds that "AI RMF core functions should be carried out in a way that reflects diverse and multidisciplinary perspectives, potentially including the views of stakeholders from outside the organization." *Id*. at 18.

and implementation[;]"[52] and (2) "MAP 1.6: Practices and personnel for design activities enable regular engagement with stakeholders, and integrate actionable user and community feedback about unanticipated negative impacts."[53] Direct community or outside stakeholder engagement is infeasible in many circumstances, since some organizations do not have the resources to work directly with outside stakeholders or are deploying low-risk AI systems that will have minimal risks. Further, the effect of AI deployment on third parties is already part of the risk assessment. Organizations should not also be required to evaluate and weigh broader "community" interests, which can often involve policy judgments that are best left to legislators and regulators.

NIST can still create alternative resources to enable impacted parties to play their "important consultative role to the primary audiences."[54] For example, NIST could make a separate document for end users, affected individuals/communities, the general public, and policy makers. The AI RMF itself, however, should be focused on individuals and entities that design, develop, and deploy AI.

B. NIST Should Directly Incorporate Existing Risk Management Tools and Standards in the AI RMF Core.

There is a significant and growing amount of work on trustworthy AI, and NIST is right to incorporate that work into its AI RMF, which is part of broader AI efforts.[55] The Second Draft layers into its discussion the OECD Recommendation on AI:2019, the OECD Framework for the Classification of AI Systems, ISO/IEC 22989:2022, ISO 31000:2018, ISO Guide 73, and ISO/IEC TS 5723:2022. NIST has also included a discussion of its work on bias and

[52] *Id*. at 20.

[53] *Id*. at 22.

[54] *Id*. at 6.

[55] *Id*. at 3 ("NIST's development of the AI RMF in collaboration with the private and public sectors is directed – and consistent with its broader AI efforts called for – by the National Artificial Intelligence Initiative Act of 2020 (P.L. 116-283), the National Security Commission on Artificial Intelligence recommendations, and the Plan for Federal Engagement in Developing Technical Standards and Related Tools. Engagement with the broad AI community during this Framework's development informs AI research and development and evaluation by NIST and others.").

explainability in AI.[56] These workstreams are relevant to the Second Draft's AI trustworthiness characteristics, and they are helpfully linked in relevant subsections. NIST's plan to create a Trustworthy and Responsible AI Resource Center to "host the AI RMF, the Playbook, and related resources to provide guidance to implement the AI RMF as well as advance trustworthy AI more broadly[]" and to accept contributions at any time is also a promising way for NIST to ensure that its AI RMF is a living document that remains relevant.[57] Still, NIST should draw on other AI resources to better align the AI RMF with existing standards work. Some examples include:

- The IEEE P7000 suite of standards,[58] especially IEEE 7000-2021, *Addressing Ethical Concerns During Systems Design*;[59]
- IEEE P3119, *Standard for the Procurement of Artificial Intelligence and Automated Decision Systems, Standard for the Procurement of Artificial Intelligence and Automated Decision Systems*;[60]
- IEEE 1012-2016, *IEEE Standard for System, Software, and Hardware Verification and Validation*;[61] and
- Publications from the ISO/IETC committee on AI, SC 42, such as ISO/IEC TF 24027:2021,[62] ISO/IEC CD 42001.2,[63] and ISO/IEC DIS 23894.[64]

[56] The Second Draft refers to NIST SP 1270, Proposal for Identifying and Managing Bias in Artificial Intelligence, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf; NISTIR 8312; and NISTIR 8367, Psychological Foundations of Explainability and Interpretability in Artificial Intelligence, https://nvlpubs.nist.gov/nistpubs/ir/2021/NIST.IR.8367.pdf.

[57] Second Draft at iii.

[58] IEEE, *Explore our approved IEEE 7000™ Standards & Projects*, https://ethicsinaction.ieee.org/p7000/.

[59] IEEE SA, *IEEE 7000-2021 Standard: Addressing Ethical Concerns During System Design*, https://engagestandards.ieee.org/ieee-7000-2021-for-systems-design-ethical-concerns.html.

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

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

[62] 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.

[63] ISO, *ISO/IEC CD 42001: 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.

[64] ISO, *ISO/IEC DIS 23894: Information technology — Artificial intelligence — Risk management*, https://www.iso.org/standard/77304.html.

NIST should also consider including informative references in the AI RMF Core. While NIST has incorporated references for particular subcategories in its Draft Playbook,[65] it should also house a core set of informative references in the AI RMF Core. This is the approach NIST has taken with the CSF. The Informative References in the CSF, along with efforts to keep them up to date with the National Online Informative References ("OLIR") Program, have proven to be efficient and effective, and NIST should follow a similar model with the AI RMF.

IV. NIST SHOULD RE-WORK THE DRAFT PLAYBOOK SO THAT IT DOES NOT UNDERMINE THE KEY FEATURES OF THE AI RMF.

A. The Playbook Could Provide Helpful Practical Guidance for Users to Apply the AI RMF if It Shares the AI RMF's Key Attributes.

CTIA supports practical tools that can help users put NIST's important risk management work into action. A companion tool like the Playbook, with appropriate adjustments discussed below, could serve this function and be a valuable addition to NIST's AI efforts.[66] As NIST explains, the Playbook "offers sample practices to be considered in carrying out th[e AI RMF] guidance, before, during, and after AI products, services, and systems are developed and deployed,"[67] and can help users "achieve the [AI RMF] outcomes through suggested tactical actions they can apply within their own context."[68] For this type of tool to be valuable, however, NIST must ensure that the Playbook is properly framed and that it maintains the same attributes of the AI RMF that NIST understands are critical to the success of risk managements tools.

[65] For each subcategory of the Playbook, NIST includes "References" that are intended to "serve as a sampling from the available literature on the given topic or subtopic area." NIST, *AI RMF Playbook FAQs*, https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook-faqs.

[66] "The Playbook provides actions Framework users could take to implement the AI RMF by incorporating trustworthiness considerations in the design, development, use, and evaluation of AI systems." *See* Second Draft at ii. NIST specifically notes that "[i]t includes example actions, references, and supplementary guidance [for the AI RMF functions]." *Id*.

[67] *Id*. at 3.

[68] *Id*. at 18.

While the Playbook should mirror *all* of the AI RMF's key attributes, CTIA sees the greatest need for improvements to the Draft Playbook with respect to the following attributes: "be risk-based, resource-efficient, pro-innovation, and voluntary;" "be useful to a wide range of perspectives, sectors, and technology domains;" "be outcome-focused and non-prescriptive;" and "be law- and regulation-agnostic."[69]

1. NIST Should Emphasize that the Playbook Is a Voluntary and Flexible Companion Tool—Not a Prescriptive Compliance Checklist.

The Draft Playbook is intended to be voluntary—just like the AI RMF.[70] The Draft Playbook is also intended to be flexible and applied to fit the fact- and context-specific needs of any given AI system, which will vary.[71] This intent is consistent with the National AI Initiative Act, which charges NIST with research and development of "best practices and voluntary standards" for AI systems and explicitly mandates that the AI RMF "not prescribe or otherwise require the use of specific information or communications technology products or services."[72]

NIST should make its discussion about the voluntary, flexible, and risk-based nature of the Playbook more explicit and noticeable—especially within the Playbook itself. In its current form, if a user were to open the Draft Playbook as a standalone tool, it could easily be mistaken for a compliance checklist because the language that promotes voluntary, flexible, and risk-based

[69] *Id*. at 4.

[70] "Like the AI RMF, the Playbook is intended for voluntary use. Organizations can utilize this information according to their needs and interests." *Id*. at ii. NIST further clarifies that "[t]he Playbook provides actions Framework users *could* take to implement the AI RMF by incorporating trustworthiness considerations in the design, development, use, and evaluation of AI systems." *Id.*

[71] NIST's Playbook FAQ explains that the suggested actions are flexible: (1) "Those interested in using the AI RMF functions to enhance their risk management posture can utilize suggested actions from the Playbook to fit their interests and needs[;]" and (2) "The NIST AI RMF Playbook is not a one-size-fits-all resource – and it is neither a checklist nor an ordered list of steps for AI actors to implement. Playbook users are not expected to review or implement all of the suggestions or to go through it as an ordered series of steps." NIST, *NIST AI RMF Playbook FAQs*, https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook-faqs.

[72] William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116-283, § 5301(c)(6), 134 Stat. 4536, 4538 (2021) ("2021 NDAA").

use of the Draft Playbook is found in external sources—either in the AI RMF or in the Playbook FAQ.

To ensure that the Playbook is clearly and explicitly voluntary, flexible, and risk-based—as is intended—NIST should consider making the following clarifications. *First,* NIST could explain on the AI RMF's "Home" page[73] that the suggested actions are not mandatory and that AI actors can implement certain actions according to their unique risk profiles and contexts. *Second*, if NIST decides to retain the presentation of the Draft Playbook—which breaks component parts into separate drop-down options and separate webpages—then NIST should include similar language about the Playbook being voluntary, flexible, and risk-based on the individual pages for each function.[74]

2. *The Draft Playbook Contains Prescriptive Language that Should Be Removed.*

In line with maintaining the AI RMF as a flexible tool that can be used universally—it is critical that the guidance provided by NIST is not prescriptive. Many actions in the Draft Playbook strike the right note. For example, many of the Draft Playbook's suggested actions lead with a verb, rather than any normative framing that an organization "must" or "should" take certain steps.[75] NIST also rightly utilizes the words "may" and "can" throughout the document.[76]

[73] NIST, *AI RMF Playbook*, https://pages.nist.gov/AIRMF/.

[74] *See, e.g.*, NIST, *AI RMF Playbook: MAP*, https://pages.nist.gov/AIRMF/map/.

[75] *See, e.g.*, NIST, *AI RMF Playbook: MAP 4.2*, https://pages.nist.gov/AIRMF/map-4/2001/04/02/risk-controls-for-third-party-risks.html ("Review third-party material (including data and models) for risks related to bias, data privacy, and security vulnerabilities.").

[76] For example, MAP 4.2's "About" section explains that "[i]n the course of their work, AI actors often utilize open-source, or otherwise freely available, third-party technologies – some of which have been reported to have privacy, bias, and security risks. Organizations *may* consider tightening up internal risk controls for these technology sources." *Id.* (emphasis added). Other Draft Playbook sections use the similarly voluntary word "can." For example, the "Transparency and Documentation" sections for each subcategory explain that "[o]rganizations *can* document the following:" and then provide example documentation practices. *See, e.g.*, NIST, *AI RMF Playbook: GOVERN 1.1*, https://pages.nist.gov/AIRMF/govern-1/2004/01/01/legal-regulatory.html (emphasis added).

In many instances, however, the language and suggested guidance in the Draft Playbook strays from this format. As an example, MAP 1.5's "About" section states, "Go/no-go decisions *should* be incorporated throughout the AI system's lifecycle."[77] Additionally, one of MAP 3.2's suggested actions explains, "[w]hen negative impacts are not direct or obvious, AI actors *should engage* with external stakeholders to investigate and document . . . "[78] This language is unnecessarily inflexible, and when it occurs amongst more voluntary "can" and "may" language, it suggests that some guidance is more rigid and prescriptive than other guidance. Accordingly, NIST should align the Playbook language with its general flexible and risk-based approach.

3. *The Playbook Should Not Reference Legal and Regulatory Regimes, Which Could Undermine the Policy-Neutral Nature of NIST's Effort.*

As detailed above, given the emerging and evolving legal and regulatory landscape with respect to AI, it is critical that the AI RMF—and any tools authored by NIST related to the AI RMF—remain policy neutral. Accordingly, while it is appropriate for NIST to note that "[l]egal and regulatory requirements involving AI are understood, managed, and documented," NIST should not attempt to provide guidance as to how an organization may satisfy such legal and regulatory requirements.[79] The Draft Playbook risks crossing this line, especially where it uses legal and regulatory guidance instead of other risk management guidance as informative references. Specifically, many of the Draft Playbook's references link to documents that could be considered to be providing legal or regulatory compliance instructions, including FTC blog posts and proposed legislation.[80] In the next iteration of the Playbook, NIST should ensure that

[77] NIST, *AI RMF Playbook: MAP 1.5*, https://pages.nist.gov/AIRMF/map-1/2001/01/05/risk-tolerance.html (emphasis added).

[78] NIST, *AI RMF Playbook: MAP 3.2*, https://pages.nist.gov/AIRMF/map-3/2001/03/02/system-cost.html (emphasis added).

[79] NIST, *AI RMF Playbook: GOVERN 1.1*, https://pages.nist.gov/AIRMF/govern-1/2004/01/01/legal-regulatory.html.

[80] *Id*. (citing to an FTC Blog Post); NIST, *AI RMF Playbook: GOVERN 2.1*, https://pages.nist.gov/AIRMF/govern-2/2004/02/01/roles-and-responsibilities.html (citing to the same FTC Blog Post); NIST, *AI RMF Playbook:*

the Playbook is law- and regulation-agnostic by not referencing existing or proposed legal or regulatory requirements. With respect to agency-specific *risk-management* tools (e.g., the Office of the Comptroller of the Currency's Model Risk Management Handbook),[81] NIST should further contextualize these informative references, noting that risk management for an institution like a bank is different from risk management for other types of organizations. As a general rule, the Playbook should only include references to non-binding academic research, industry standards, or industry or government risk-management tools—not to laws or regulations.

Additionally, as with the AI RMF, the Playbook should be explicit that not all trustworthiness characteristics will be relevant to all AI systems.

B. The Playbook Should Maximize the Practical Utility and Usability of the AI RMF.

1. The Playbook Should Direct Guidance Towards Individuals and Enterprises that Design, Develop, and Deploy AI—Not Towards Individual End Users or the Broader Community.

Like the AI RMF, in numerous parts of the Draft Playbook, NIST suggests that organizations should create processes for engagement with stakeholders and incorporate their feedback during AI system design and implementation.[82] As discussed above, this expectation is infeasible because risk profiles and AI deployment contexts vary. Organizational resources for

GOVERN 4.2, https://pages.nist.gov/AIRMF/govern-4/2004/04/02/team-empowerment.html (citing to the Algorithmic Accountability Act of 2019 (H.R. 2231, 116th Cong. (2019)).

[81] *See, e.g.*, NIST, *AI RMF Playbook: MAP 4.1*, https://pages.nist.gov/AIRMF/map-4/2001/04/01/document-third-party-risks.html.

[82] *See* NIST, *AI RMF Playbook: GOVERN 1.2*, https://pages.nist.gov/AIRMF/govern-1/2004/01/02/trustworthiness-embedded-into-processes.html ("Outline processes for internal and external stakeholder engagement."); NIST, *AI RMF Playbook: GOVERN 5.2*, https://pages.nist.gov/AIRMF/govern-5/2004/05/02/post-feedback.html ("Mechanisms are established to enable AI actors to regularly incorporate adjudicated stakeholder feedback into system design and implementation."); NIST, *AI RMF Playbook: MAP 1.6*, https://pages.nist.gov/AIRMF/map-1/2001/01/06/stakeholders.html ("Practices and personnel for design activities enable regular engagement with stakeholders, and integrate actionable user and community feedback about unanticipated negative impacts."); NIST, *AI RMF Playbook: GOVERN 3.1*, https://pages.nist.gov/AIRMF/govern-3/2004/03/01/team-composition.html ("Stakeholder involvement: Include diverse perspectives from a community of stakeholders throughout the AI life cycle to mitigate risks."); NIST, *AI RMF Playbook: GOVERN 5.1*, https://pages.nist.gov/AIRMF/govern-5/2004/05/01/capturing-stakeholder-feedback.html.

outside stakeholder engagement may be slim for smaller or younger organizations. Additionally, some AI deployment contexts will have very few risks. NIST should tailor the Playbook's audience to AI system designers, developers, and deployers who will benefit from a clear tool regarding the AI risk management process, rather than to a broader set of stakeholders.

2. *The Playbook Should Be Used as a Tool to Show the Range of Ways the AI RMF May Be Applied Across Variable Use Cases.*

NIST should adjust the Draft Playbook's subcategories to better account for the wide range of organizations that will be using the AI RMF, and the broad range of use cases to which the AI RMF will apply. In some subcategories, the Draft Playbook presupposes that large organizations are the intended users. For example, MAP 1.2 recommends that organizations, "Create and empower interdisciplinary expert teams to capture, learn, and engage the interdependencies of deployed AI systems and related terminologies and concepts from disciplines outside of AI practice such as law, sociology, psychology, anthropology, public policy, systems design, and engineering."[83] This is a highly aspirational recommendation that will be infeasible for many organizations to achieve.

While the Draft Playbook makes a helpful distinction between high- and low-risk AI systems in some places,[84] this distinction and theme should carry through the document. NIST can go further and provide more descriptions of low-risk systems to increase the Playbook's practicality. For example, NIST should give specific examples of low-risk systems in the "About" section for GOVERN 5.2.[85] As explained above, separate, flexible guidance for low-risk systems will allow users to conserve AI-related resources. A more detailed description of

[83] NIST, *AI RMF Playbook: MAP 1.2*, https://pages.nist.gov/AIRMF/map-1/2001/01/02/ai-actors.html.

[84] For example, GOVERN 5.2 states, "Organizations should apply a risk tolerance approach where higher risk systems receive larger allocations of risk management resources and lower risk systems receive less resources." NIST, *AI RMF Playbook: GOVERN 5.2*, https://pages.nist.gov/AIRMF/govern-5/2004/05/02/post-feedback.html.

[85] *Id*.

low-risk systems will allow users to navigate the Playbook more efficiently and effectively. To account for various AI developers/deployers as well as a wide array of use cases, NIST should rework subcategory descriptions as well as their corresponding suggested actions to be more generally applicable to significantly more AI systems and their owners/operators.

V. NIST SHOULD ENSURE THAT ADVANCEMENTS AND ADDITIONS TO THE AI RMF ARE ITERATIVE AND BUILD ON SOLID FOUNDATIONS.

A. NIST Should Finalize the AI RMF and the Playbook Before Starting Work on the AI RMF Profiles.

In the Second Draft, NIST explains that it is soliciting contributions towards the development of AI Profiles. While requesting contributions for possible profiles to be included in the NIST Trustworthy and Responsible AI Resource Center may aid NIST and various stakeholders, NIST should refrain from encouraging the development of AI Profiles until the AI RMF and the Playbook are finalized in 2023.

CTIA supports the development of profiles, which have been successfully used in deploying the CSF, resulting in such guidance as the Positioning, Navigation, and Timing ("PNT") Profile and the Ransomware Profile.[86] However, delaying any substantive work on profiles until the first versions of the AI RMF and Playbook are complete will to ensure that profiles have the full benefit of public collaboration and are based on the final versions of NIST's risk management tools. Indeed, there have been significant changes from the Initial Draft to the Second Draft. If there are similar changes from the Second Draft to the final version, a profile based on the Second Draft would neither be helpful nor relevant.

[86] *See* NIST, *Responsible Use of Positioning, Navigation and Timing Services*, https://www.nist.gov/pnt (last visited Sept. 20, 2022); NIST, NISTIR 8374, Ransomware Risk Management: A Cybersecurity Framework Profile (Feb. 2022), https://doi.org/10.6028/NIST.IR.8374.

B. NIST Should Clarify the Playbook's Update Schedule and Better Align It with the AI RMF's Schedule.

In the Playbook's FAQ, NIST explains that the AI RMF will be updated "from time-to-time" while the Playbook will be "more dynamic" and likely updated more frequently with "no final version."[87] Additionally, comments on the Playbook can be suggested at any time and "will be reviewed and integrated on a semi-annual basis."[88]

While CTIA supports the living nature of the document and supports robust and regular collaboration between NIST and various stakeholders, NIST should better solidify the process for updating the Playbook and harmonize the Playbook's update schedule with the AI RMF's schedule. Specifically, NIST should publish a "final" Playbook, just as it intends to publish a "final" AI RMF. Refraining from finalizing the Playbook will lead to confusion and will undercut the practical value of the tool. A document like the Playbook can be both a living document with periodic revision, as well as complete at certain points in time.

NIST should use the same comment and review process for Playbook updates that it uses for the AI RMF, which would be more efficient and allow NIST to receive more consistent input from the AI community. At a minimum, NIST should provide a significant amount of time to review and comment on the revised Playbook sections as well as the other half of the Playbook when both are released.

VI. CONCLUSION.

CTIA appreciates the opportunity to provide feedback to NIST and looks forward to continued collaboration.

[87] NIST, *NIST AI RMF Playbook FAQs*, https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook-faqs.
[88] *Id.*

Respectfully submitted,

/s/ Avonne S. Bell
Avonne S. Bell
Director, Connected Life

Thomas K. Sawanobori
Senior Vice President and Chief Technology Officer

John A. Marinho
Vice President, Technology and Cybersecurity

Justin C. Perkins
Manager, Cybersecurity and Policy

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

September 29, 2022