August 29, 2022
Dr. Laurie E. Locascio
Director
National Institute of Standards and Technology (NIST)
Re: Artificial Intelligence Risk Management Framework Second Draft

Dear Director Locascio,

Accenture is a global professional services company with leading capabilities in digital, cloud and security. Combining unmatched experience and specialized skills across more than 40 industries, we offer Strategy and Consulting, Technology and Operations services and Accenture Song - all powered by the world's largest network of Advanced Technology and Intelligent Operations centers. Our 710,000 people deliver on the promise of technology and human ingenuity every day, serving clients in more than 120 countries. We embrace the power of change to create value and shared success for our clients, people, shareholders, partners and communities.

Accenture appreciates the opportunity to provide input on the second draft of NIST's AI Risk Management Framework, and the first draft of NIST's RMF Playbook. We commend NIST for its longstanding and ongoing facilitation of public private partnerships to develop such useful devices as the Cybersecurity and Privacy Risk Management Frameworks. These inclusive processes and the frameworks they produce help educate stakeholders and provide actionable iterative guidance and best practices to manage risks of using technology.

The comments that follow build upon our past support for the NIST RMF process and intend to maximize its impact and ensure its long-term relevance. We look forward to further participation and collaboration with NIST on this and future initiatives.

Sincerely,

Paul Daugherty

Chief Technology Officer

Accenture

Introduction

The NIST AI RMF and RMF Playbook appear well-positioned to provide actionable guidance to organizations managing risks related to the development and use of AI. Accenture's comments highlight areas that could benefit from increased clarification, amendment, or further consideration.

While the RMF provides an excellent step forward, most companies (69%) have started implementing Responsible AI practices, but only 6% have operationalized their capabilities to be responsible by design. Being responsible and managing risk by design will become more beneficial over time, especially as governments and regulators consider new standards for the development and use of AI. Countries such as the United Kingdom, Brazil, and China are already acting, either by evolving existing requirements related to AI, or through the development of new regulatory policy.[1]

The ability to deliver and implement high quality, trustworthy AI systems will increase trust among AI consumers and offer first movers a significant advantage in the short-term, enabling them to attract new customers, retain existing ones, and build investor confidence. Once published, NIST and the Department of Commerce should find ways to promote the RMF and Playbook, and Accenture stands ready to support such efforts. Governments will eventually require organizations to do many of the things illustrated by the RMF and the Playbook. Therefore, rapid and robust consideration of these materials is in our commercial interests as we serve the interests of our customers and society.

As the RMF develops, we also encourage NIST to draw attention to the key commonalities and differences between the RMF and emerging global thinking on this topic (e.g., the draft EU AI Act). Accenture and its clients are operating across geographical boundaries and seeking to promote and implement the responsible use of data and AI. The RMF has a window of opportunity to lead the approach organizations take, which will only be enhanced further if NIST was to demonstrate how it supports organizations operating on a global scale.[2]

The below directly responds to NIST's questions about the RMF and the Playbook.

Does the AI RMF enable decisions about how an organization can increase understanding of, communication about, and efforts to manage AI risks?

1. The RMF will help organizations understand, communicate, and manage AI risks. We continue to be encouraged that both the initial and most recent draft are directly in-line with Accenture's approach to Responsible AI. We also applaud NIST's inclusion of our suggestion to create a companion RMF Playbook.

   In its current form, the RMF is a useful document for technologists, IT professionals and the

[1] From AI Compliance to Competitive Advantage https://www.accenture.com/us-en/insights/artificial-intelligence/ai-compliance-competitive-advantage

[2] GDPR – An Opportunity in Disguise https://www.accenture.com/_acnmedia/pdf-83/accenture-gdpr-opportunity-disguise.pdf

other actors (page 6) highlighted by OECD's Framework for the Classification of AI systems.[3] However, the RMF also points out that 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." In practice, this means C-suite leaders, and more needs to be done to reach that audience.

We strongly recommend that, as part of its publication and eventual promotion of the NIST AI RMF and Playbook, NIST leverage the Business Roundtable Roadmap for Responsible AI and consider partnering with them on events. The Roadmap reflects the collective views of leading CEOs representing every aspect of the complex AI ecosystem — including some of the world's largest developers and users of AI as well as companies that are just beginning their AI journey. Given the current state of regulation and market forces, CEO leadership is needed to embed these practices throughout an organization, as well as ensure that future AI developments and deployments are designed responsibly.[4]

As mentioned previously, the RMF may benefit with a side-by-side comparison of the RMF, draft EU AI Act and other emerging regulatory frameworks. The investments required to comply with the eventual EU AI Act, for example, are expected to be substantial, and similar to what multi-national companies faced with the General Data Protection Regulation, key investment decisions will need to be made over the course of the next two years around data infrastructure, AI governance and internal changes to policies and processes.

2. Accenture believes the AI RMF Second Draft is flexible enough to serve as a continuing resource considering the evolving AI technology and standards landscape, and strongly support NIST's efforts to ensure that both the AI RMF and RMF Playbook continue to be flexible, living documents. Moving forward, as new technology emerges, we expect developments will require iterative updates and more catered risk profiles. For example:
    a. Neuromorphic Computing technology will enable "intelligence at the edge" allowing AI systems to operate without a connection to central repositories of information or centralized control and oversight. While the overall risk concerns around AI on neuromorphic chips are largely in line with those addressed in the framework, the scale of deployment is likely to be much larger, and this technology may enable deployment of AI systems in locations that are difficult to access for updates, reviews, and tests of AI decision-making and other processes. In this case, the suggestions made in the RMF are flexible enough to continue to apply, but they will be even *more* important given the possible difficulty of addressing issues post-deployment.
    b. Breakthroughs in Natural Language Processing Large Language Models (LLMs), as well as text-to-image generators such as Dall-E,[5] are creating more natural AI generated content. These tools will enrich people's lives but also present challenges for enterprises

[3] OECD Framework for the Classification of AI systems https://www.oecd-ilibrary.org/science-and-technology/oecd-framework-for-the-classification-of-ai-systems_cb6d9eca-en
[4] Business Roundtable Roadmap for Responsible AI https://www.businessroundtable.org/policy-perspectives/technology/ai
[5] DALL·E 2 (openai.com)

seeking to leverage them in determining the risk and accountability mechanisms needed to manage their use. The RMF provides a structure for such an assessment but may need to consider such novel developments on a case-by-case basis and provide supplementary guidance to ensure organizations can apply the RMF appropriately, confidently, and practically.

Does the AI RMF appropriately cover and address AI risks?

1. The RMF lays out strong and appropriate considerations for analyzing, mitigating, and creating clear guidance for minimizing potential harms and risks related to AI. We believe one key piece is missing.

   In many cases, AI is augmenting or replacing what was once a primarily human activity. In those cases, effective risk management requires that potential risks be considered in a framework that includes a comparison to those alternatives (i.e., how is the risk-profile changing by moving from a human activity, to one performed by AI). This continues to be important as the risk profile assessment for an AI system will need to take into account the typical harms or risks associated with the way the action/process might have operated when it was manually performed versus when it is performed by the AI system versus when the context of operation is a hybrid between human and AI. NIST should consider further developing the RMF to help organizations understand how to identify and manage AI related risks within the different socio-political and technical contexts of use, and how to evaluate the risks associated with the use of AI vs risks of not using AI.

   Is there anything else missing from the AI RMF Draft?

1. As AI technologies continue to be adopted across industries, Accenture supports NIST's use of use-case profiles based on the requirements and risk tolerance required in those cases. We suggest NIST consider developing use-case profiles beginning with:
   - Human resources and talent management
   - Health
   - Public health services
   - Life sciences
   - Synthetic drug development
   - Lending & credit
   - Content moderation

2. Accenture also suggests that NIST monitor other legislative and regulatory developments such as the draft EU AI Act, in order to provide guidance to industry on how to apply the RMF to high-risk and potentially high-risk use cases. This may not apply to the version of the RMF that will be published in January but will be critical as NIST considers updates after regulations are enacted.

3. It would be beneficial if the AI RMF could provide guidance to help organizations anticipate and manage the risks of buying, supplying and using AI through the lifecycle (in the same way as we might approach managing product liability risk in manufacturing, for example).

RMF Playbook

1. We welcome NIST's creation of the accompanying Playbook. While incomplete, it appears to have clear guidance and controls that an entity will want to consider when developing or deploying AI.
    - NIST has intimated that the presentation mode and style for the Playbook will be considered separately. We welcome this, as different audiences will engage with the Playbook for different reasons, and through different mechanisms. While the Playbook draft is well suited for technologists' use, we recommend creating different guidance for other key audiences, e.g. legal experts and C-suite leaders, who have different objectives and interests (e.g., governance).
    - Building upon that, NIST should also consider various events, engagements, and exercises that could engage C-suite leaders in necessary conversations to implement the suggestions in the "govern" section. We envision this as an ever-changing promotional campaign—with the core tenets of the Playbook draft being the underlying resource.

2. In the Map function, section 1.1, NIST suggests that "[organizations should] pursue AI system design purposefully, after non-AI solutions are considered." This language can be read to imply that organizations should not consider an AI solution until all other possible approaches are explored and ruled out. We think this recommendation may unduly encourage organizations to seek manual rather than automated solutions in all situations and is likely not what NIST intended with this suggested action.