From: Andrew Clark
Sent: Thursday, September 29, 2022 4:09:32 PM (UTC-05:00) Eastern Time (US & Canada)
To: aiframework <aiframework@nist.gov>
Subject: NIST AI Playbook Feedback

Dear NIST AI Playbook team,

I think the idea of having a playbook as a supplement to the AI Risk Management Framework is fantastic. Having an easy-to-implement playbook to help organizations with the daunting task of going from nothing to something is a very much-needed tool. However, in the Playbook's current form, it is not as intuitive and easy to use as it could be. Below I have enumerated several areas for improvement. Monitaur is willing to assist in making these changes to the framework and playbook:

- Map categories into a spreadsheet that can be tied back to existing risk and control structures already in use by organizations. The NIST Cybersecurity framework did an excellent job of this, even going as far as linking to ITIL, ISO, and COBIT controls. For a framework to be useful to organizations it needs to be in a risk and controls framework that is mapped to their existing risk management processes.
- The Playbook has excellent points and thoughts but at times, it is duplicative. As an example, Govern 6.1 and 6.2 are, in practice, the same as Map 4.1 and 4.2. The risk is that adequate controls are in place around third parties. Once those risks are defined, from a practitioner's perspective, we put mitigating controls in place to ensure the risks are mitigated. Blurring the respective 3 lines of defense roles, once the controls are defined, they would be implemented and periodically assessed by both internal and external parties for effectiveness. Creating a 'risk universe' and 1-to-many controls to mitigate each risk is how many organizations approach these problems.
- The Playbook could be implemented more effectively, potentially framed as a "simplified matrix" for early adoption and an "advanced matrix" for meaningful long-term integration, as touched on in the previous point. In its current form, it needs to be mapped into a risk and controls structure to be readily actionable.
- The structure of the individual 'controls' is repetitive and occasionally confusing. Instead of having an implied control title with: 'About', 'Actions', and 'Transparency and Documentation' sections, move to a standard structure of the form: 'Control Title', 'Description', and 'Evidence Documentation', and 'Motivating risks' would allow better integration into existing risk management frameworks. Additionally, the 'About' section should be broken down further and any identified risks should be mapped into a 'Risk Title', 'Description', and 'Mitigating Control[s]' structure.
- When an auditor or risk management would ascertain if a model is being governed properly, they will send out a request for information. Each piece of the requested information should be included in one control and asked for once. In the current structure, documentation is requested multiple times. As an example, Govern 3.1 asks for documentation on "Stakeholder involvement: Include diverse perspectives from a community of stakeholders throughout the AI life cycle to mitigate risks.". In Map 1.2, it asks for documentation on: "To what extent do the teams responsible for developing and maintaining the AI system reflect diverse opinions, backgrounds, experiences, and perspectives?", etc. The four documentation items could be addressed in the format: "Risk: Lack of diversity and interdisciplinary perspectives on the creation, management, and assessment of the AI System" with 1-to-N controls outlining interdisciplinary team requirements, periodic independent evaluation of system performance, and impact, etc. Approaching governance in a one-to-many relationship of one risk to potentially multiple, unique, non-overlapping controls provides an actionable, unambiguous approach that integrates into existing risk management practices. A common best practice that

many organizations are employing is called the “common controls framework”. As many organizations have overlapping compliance considerations, organizing their controls into ‘one pane of glass, creates efficiency savings once the risks and controls are mapped, compliance tasks and documentation need to only be performed once. An excellent, succinct, example of this approach can be seen by Atlassian here: Atlassian's Common Controls Framework | Atlassian

The NIST AI Playbook authors have put in a lot of work to get a solid first draft. However, before it can be useful to organizations, many changes will be required, such as making it less duplicative and structuring it in a way that is conducive to existing risk management processes. The more seamlessly NIST’s AI RMF efforts can fit into these existing frameworks and approaches without causing ambiguity and duplication of work, the wider its adoption will be.

Sincerely,

Dr. Andrew Clark

Additional information:
I am an experienced auditor and ML audit expert with years of experience as an ISACA ML Assurance SME presenting. I’ve presented 5 times at their North America conference on ML Assurance; have been an expert reviewer on their Auditing AI guidance, and published on ML auditing in their journal. I would be honored to help with future versions of the playbook. I am happy to send across supporting resources, papers, and presentations if desired.

My bio is below:
Dr. Andrew Clark is Monitaur’s co-founder and Chief Technology Officer. A trusted domain expert on the topic of ML auditing and assurance, Andrew built and deployed ML auditing solutions at Capital One. He has contributed to ML auditing education and standards at organizations including ISACA and ICO in the UK. Before Monitaur, he also served as an economist and modeling advisor for several very prominent crypto-economic projects while at Block Science.
Andrew received a B.S. in Business Administration with a concentration in Accounting, Summa Cum Laude, from the University of Tennessee at Chattanooga, an M.S. in Data Science from Southern Methodist University, and a Ph.D. in Economics from the University of Reading. He also holds the Certified Analytics Professional and American Statistical Association Graduate Statistician certifications.