TECHNET
THE VOICE OF THE
INNOVATION ECONOMY

1420 New York Avenue NW, Suite 825
Washington, D.C. 20005
www.technet.org | @TechNetUpdate

May 31, 2024

National Institute of Standards and Technology
100 Bureau Drive, Mail Stop 8900
Gaithersburg, MD 20899

RE: NIST AI 600-1, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile

To Whom It May Concern:

TechNet appreciates the opportunity to comment on the National Institute of Standards and Technology's (NIST) companion resource for Generative AI to the AI Risk Management Framework (NIST AI 600-1). Many of our nation's leading AI developers, deployers, researchers, and users are TechNet members.

TechNet is the national, bipartisan network of technology CEOs and senior executives that promotes the growth of the innovation economy by advocating a targeted policy agenda at the federal and 50-state level. TechNet's diverse membership includes dynamic American businesses ranging from startups to the most iconic companies on the planet and represents over 4.4 million employees and countless customers in the fields of information technology, artificial intelligence, e-commerce, the sharing and gig economies, advanced energy, transportation, cybersecurity, venture capital, and finance.

Unique Roles in the AI Ecosystem

We appreciate NIST AI 600-1's acknowledgment of AI actors' unique roles. We believe it is crucial that, as these policies evolve, policymakers should be clear that the responsibilities of developers, deployers, and users must align with the unique roles each of them play throughout the AI life cycle. While everyone is responsible for ensuring AI is trustworthy, every actor in the ecosystem will serve a different role in safety implementation. In order to implement effective safety policies, it is necessary for all of these actors to work together and have clear roles in maintaining responsible AI systems.

We would appreciate additional guidance from NIST for the open-source ecosystem. This profile recommends some actions that may be infeasible for open-source developers, such as overseeing downstream deployers and decommissioning Generative AI systems. As additional open-source Generative AI tools come into the market, tailored guidance from NIST would be beneficial to support responsible deployment.

Feedback Requirements

The AI startup ecosystem is vital to maintaining America's competitive edge in the global economy. We would appreciate NIST's consideration of the potential recommendations for small and mid-size businesses. We are concerned that the many feedback recommendations in the draft profile may be impractical or overly burdensome for smaller organizations with limited resources. We agree with NIST that external feedback, such as information gained during red-teaming, is beneficial for model performance and safety, but we would appreciate ranked guidance amongst the many processes recommended for third-party input (such as GV-1.5-007, GV-3.2-005, GV-5.1-003, MP-5.1-009, MS-2.8-014, MS-3.3-009, MS 4.2-003).

Training Data Requirements

TechNet believes that the draft profile's restrictions on the use of copyrighted data should be revisited. We are concerned that these recommendations are based on policy preferences and lack the scientific approach needed to allow American AI developers and deployers to continue innovating and building next-generation tools. The goal of Generative AI is to help authors create new content — whether it consists of responses to user queries, longer pieces of text, visual works, music, or computer code. Encouraging the creation and dissemination of new expression is the very purpose of copyright law. TechNet believes that existing copyright doctrines are sufficiently flexible to handle many of the scenarios that will likely arise with Generative AI and that courts — informed with the facts of specific cases — are the appropriate first venues for determining how those doctrines should apply. We do not believe it is appropriate for NIST to state that there is an inherent risk in the use of copyrighted materials, as other U.S. government agencies and courts are currently examining this question.

NIST AI 600-1 includes recommendations that private data should be anonymized when used in training data or that organizations need individual consent to use personal data for training purposes. We agree that organizations should employ best practices to anonymize personal data wherever feasible and appropriate. Our concern is that some of these considerations may not be feasible to comply with in practice. For example, it may be beneficial to the testing of a system to process personal data to ensure it operates as intended, such as in bias testing or other evaluations. We believe NIST AI 600-1's recommendations should look to balance privacy protection with the practicalities of AI training.

Given the size of training data for Generative AI models, it is difficult for organizations to ensure that harmful and illegal content is not included in these data sets. TechNet members undertake best practices to remove this data but cannot universally guarantee that all targeted content will be removed. This challenge is acknowledged in section 5.1.1 of NIST's guidance on Reducing Risks Posed by Synthetic Content (NIST AI 100-4). We recommend that action GV-1.2-005 be rephrased to account for this. This draft profile would also benefit from greater clarity on what types of sensitive information should be removed from training data to manage CBRNE concerns in a way that does not inordinately hamper the model's ability to provide scientific outputs. We are also

concerned about the conflation of illegal and harmful content and urge these categories to be separated, especially as "harmful" is often context-dependent. We must also point out that potentially harmful content can be useful in training models to detect and filter out other harmful content. We want to ensure that America's AI developers have the tools necessary to build responsible systems and not inadvertently limit safety practices.

Provenance Data Tracking

TechNet is concerned that NIST AI 600-1 conflicts with the provenance policies put forth in NIST AI 100-4 (Reducing Risks Posed by Synthetic Content). While NIST AI 100-4 looks to outline best practices and current research on provenance data tracking, NIST AI 600-1 moves ahead and establishes new processes and reporting requirements for provenance. We are concerned that this will lock in a process that isn't built upon the most effective provenance data tracking methods, as this important research and standardization is still ongoing. We advise that NIST AI 600-1 incorporate or direct to the processes underway in the development of NIST AI 100-4.

Independent Assessments

We appreciate the NIST AI Safety Institute's efforts to develop guidance and benchmarks for evaluating AI capabilities and believe this effort can build upon existing industry processes. However, TechNet believes that NIST AI 600-1's recommendations to procure independent third-party auditing is premature. Mandating an independent audit before appropriate technical standards and conformity assessment requirements are established could open Generative AI systems to trade secrets theft and inaccurate audit reports. We would also appreciate greater clarity from NIST on what types of data should be subject to audits. We believe these data audit recommendations should be carefully developed to avoid creating new risks or leading to the overexposure of training data. We look forward to partnering with NIST and other evaluation organizations to develop additional research on the best way to assess AI systems.

Conclusion

We look forward to working with you on AI policy and appreciate the opportunity to discuss this innovative technology. We stand ready to serve as a resource to you in your examination of this important issue. Thank you for your consideration of our perspective.

Sincerely,

Linda Moore
President and CEO