

Commentary for RFI Related to NIST's Assignments Concerning Artificial Intelligence

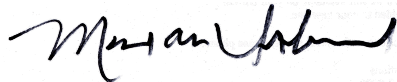
Under Sections 4.1, 4.5 and 11 of the Executive Order

2 February 2024

Rachel Trello
NIST AI Guidance Team NIST-2023-0309
100 Bureau Drive, Stop 8900, Gaithersburg, MD 20899
ai-inquiries@nist.gov

Please find attached my comments regarding the NIST AI Executive Order.¹

Regards,



Mark Alan Underwood CIRM, CDPSE, CSQE
Krypton Brothers LLC
7 Birchwood Ave
Port Washington NY 11050
516.234.0076 (m)
mark.underwood@kryptonbrothers.com

¹ RFI solicitation <https://www.regulations.gov/document/NIST-2023-0009-0001>

1. **Liaison with IEEE standards and working groups.** Beginning with IEEE 7000 ([IEEE Standard Model Process for Addressing Ethical Concerns during System Design](#)), IEEE working groups have considered various facets of AI. Newer working groups, including IEEE P2957 [Standard for a Reference Architecture for Big Data Governance and Metadata Management](#) and IEEE P3396 [Recommended Practice for Defining and Evaluating Artificial Intelligence \(AI\) Risk, Safety, Trustworthiness, and Responsibility](#) (to name two among several) are developing related frameworks.
2. **IEEE P2957** I chair the IEEE 2957 Working Group, and see a considerable importance for metadata as it relates to data and model provenance for AI.² Metadata helps to address a principal concern from the public: “Is this picture an AI?” This is a major concern of musicians, as seen in [this Credits Due proposal from Abba cofounder Bjorn Ulveus](#) which aims “to tackle the music industry’s royalty payment and *metadata* problem.”
3. **Metadata and ABAC** To the extent that “attributes” and “metadata” are essentially interchangeable, the IEEE P2957 effort is connected to the NIST SP 800-162 Guide to Attribute Based Access Control (ABAC) Definition and Considerations. This connection has implications for security, assurance and trust for model-based AI, such as LLM’s, and for “traditional” formal AI methods, such as ontologies. I note that the proposed framework efforts in response to the Executive Order (EO) will leverage much of NIST’s existing work in security, privacy and risk. This is commendable. However, SP 800-162 is not adequately integrated into other NIST standards; it is especially needed for 800-53 and 800-37 (the RMF) to support controls and capabilities such as observability and provenance. It is also needed to track as-used instances of AI applications (“use cases,” or instances) both at a point in time and longitudinally as an application (or user sophistication) matures.
4. **SBOMs** The AI framework should be coordinated with work on SBOMs. Existing notions of software services can readily be extended to include AI managed components.
5. **SDLC Reconsideration** NIST documents and standards related to the software development life cycle (SDLC) should be revisited in light of the possibility that products, tools and processes may be AI (code generated through automated means).
6. **NIST OSCAL** The NIST OSCAL model should be extended to include assurance and conformance monitoring for AI elements.
7. **NIST DIOPTRA** I support the parallel development of DIOPTRA to improve the usability and feasibility of this work.

² See Underwood, M. (2023). Continuous Metadata in Continuous Integration, Stream Processing and Enterprise DataOps. *Data Intelligence*, 5(1), 275–288. https://doi.org/10.1162/dint_a_00193

8. **Intelligent Agents** In an earlier era of NIST and AI's development, *intelligent agents* received attention. It may prove fruitful to revisit this earlier work now that agents can be easily and ubiquitously integrated in human-machine systems. AI systems will likely include multiple AI and human agents with new and emerging capabilities that were not envisioned in existing NIST documents. One suspects that most of the NIST work in this area simply needs augmentation for these now more salient use cases, but in some cases this may need some rethinking – e.g.; the very notion of “system design” may become more cut-and-paste (LLM-driven) than based on the drawing-board / blueprint model. Red-teaming, for instance, can be seen as a multi-agent, cyber-range based framework with both human and AI agents taking multiple roles based on risk, domain-specific consideration (human health and safety vs. systems safety) and the priority given to mitigation and resilience.
9. **Domain-specific AI Engineering** You ask about “the types of professions, skills, and disciplinary expertise organizations need to effectively govern generative AI, and what roles individuals bringing such knowledge could serve.” As noted in the [NIST Big Data Interoperability Framework](#) 1500, Volume 3 Security and Privacy, a more prominent role is needed for domain specialists. This is a deep topic which can't be covered thoroughly here, but is addressed in some detail there. In a nutshell: security and privacy for radiology AI is substantially different from security and privacy for autonomous vehicle AI. There are important distinctions that cannot be left to computer scientists, data scientists or developers – nor to AI itself.
10. **Assurance and Audit** AI governance calls not only for an integration of AI processes to assist with conformance and data gathering. It also needs assurance processes, which tend to be left to periodic audits and manual processes. Existing drafts I've studied haven't fully embraced “RegOps” / “ModelOps” concepts, but a newer AI framework could.
11. **Driving adoption and implementation of AI-related standards** While standards are tool- and product-agnostic (and silent), the standards need to be able to “work with” products which have already gained market share. In today's world, that means a standard that was unusable with OpenAI's ChatGPT API's would create a hurdle to adoption. An open source equivalent, where available, is a preferable alternative – if available.
12. **LLM Dog Food** Also, the standards themselves need to be tested with LLM ingest to understand their outputs and suitability. More focus on practical use cases, indeed, a cut-and-paste composable framework needs to be made harmonious with software engineering best practices. Put the standards into an LLM. See if the generated guidance for a specific domain is seen by ethics and domain specialists as compatible with the objectives of the standard.

13. **Adopt Persistent Use Cases** Because domain specificity matters, use cases should be incorporated across the various NIST security, privacy and AI domains. Too often, each document stands up its own use cases, which inhibits cross-adoption. Perhaps these use cases need persistent web identifiers a la linked web, so that invoking them in standards language is facilitated.