202 Burlington Road
M/S M300
Bedford, MA 01730

January 27, 2019

Submitted Electronically
Katie MacFarland
National Institute of Standards and Technology
100 Bureau Drive
Stop 2000
Gaithersburg, MD 20899

Dear Ms. MacFarland:

Thank you for the opportunity NIST has afforded the public to provide input on Developing a Privacy Framework via its November 14, 2018 Request for Information (Docket Number 181101997-8997-01). In response, we provide the following comments in our individual capacities as practicing privacy engineers.

3. How organizations define and assess risk generally, and privacy risk specifically.

The continued evolution of privacy, data governance, and enterprise risk management (ERM) means that any attempt to gauge how organizations define and assess general and privacy-specific risk is taking aim at a moving target. However, irrespective of the details at any of these broad levels, a key issue for many organizations is the interfaces between them. Beyond the standard issue of reconciling qualitative and quantitative approaches to risk, there is the question of how to "roll up" fundamentally incommensurate types of risk. Even absent any attempt to actually aggregate risk determinations, how privacy risk is understood and compared with other types of risk remains a fundamental problem. How, for example, should privacy risk be traded off against security risk in terms of data governance or against safety risk in terms of ERM? Operating effectively within the applicable risk trade-spaces is one of the primary goals of data governance and ERM, but this presumes the necessary normalization mechanisms.

Attempts to render privacy risk more quantitative have generally produced numbers lacking any grounded meaning. This typically takes the form of judgements made on an ordinal scale with no defined basis for differentiating between one value and another. In some cases, there is not even a clear definition of the metric itself. As such, these numbers convey a false sense of precision. More systematic quantitative approaches to privacy risk are possible, but require significant data together with a willingness to reason from that data and to invest in calibration exercises for consistency.

Qualitative approaches to privacy risk (historically in the form of privacy impact assessments) can be far more useful than arbitrary quantitative approaches, but suffer both from huge variations in analytical quality and the constraints of narrow privacy risk models (FIPPs in particular).

Given this, we believe that the Privacy Framework should be agnostic with respect to quantitative versus qualitative bases for assessing privacy risk, but should emphasize or encourage the use of formal or rigorous methods for doing so and define the qualities these terms represent. The Framework could add particular value to the extent it can facilitate normalization across different levels of risk management—privacy, data governance, and ERM—and between privacy and other types of risk.

10. What standards, frameworks, models, methodologies, tools, guidelines and best practices, and principles organizations are aware of or using to identify, assess, manage, and communicate privacy risk at the management, operational, and technical levels…

The PIA and its GDPR-inspired cousin, the data protection impact assessment (DPIA), continue to serve as the principal approach to privacy risk of all types in both the private and public spheres. Part of their appeal, undoubtedly, is their typically all-in-one nature, attempting to address at management, operational, and technical levels the identification, assessment, management, and communication of privacy risk. This monolithic aspect is also seen in other, more ambitious methodologies, such as the OASIS Privacy Reference Management Model and Methodology (PMRM) [1] and the PRIPARE Privacy and Security-by-design Methodology[2].

While appealing, monolithic approaches are problematic by definition as they are designed to be literally all or nothing. This, by itself, undermines several of the desired attributes of the Privacy Framework—adaptability, integrability with ERM processes, and compatibility with other approaches (which may themselves be monolithic). Ideally, approaches (whatever specific form they take or label that is applied to them) should be sufficiently modular that they can be mixed and matched to address the different levels and activities of privacy risk management.

Privacy risk models (typically characterizing one or more of threats, vulnerabilities, and consequences) are one of the most obvious examples of this. At this point there are multiple privacy risk models, exhibiting various degrees of completeness, that can be used to identify privacy risk. These include, but are not limited to, FIPPs, Solove's taxonomy[3], Nissenbaum's contextual integrity[4], and the one described in NISTIR 8062 (and in more detail in its initial draft). Ideally, any method for assessing privacy risk should be capable of using the privacy risk model of choice, as does, for example, System-Theoretic Process Analysis for privacy (STPA-Priv)[5]. Similarly, it should be possible to utilize distinct approaches to, for example, privacy risk at the operational and technical levels. The point is not that this is necessarily desirable in every instance, but that it should be enabled to the maximum extent practicable. One way of achieving this would be through use of a metamodel, as discussed further in our response to Question 18.

---

[1] Organization for the Advancement of Structured Information Standards (OASIS), "Privacy Management Reference Model and Methodology (PMRM) Version 1.0," Committee Specification 02, Burlington, MA: OASIS, 2016.
[2] PReparing Industry to Privacy-by-design by supporting its Application in REsearch (PRIPARE), "Updated Privacy and Security-by-design Methodology" (Deliverable D1.3, 24 September 2015), Paris, 2015.
[3] D. Solove, *Understanding Privacy,* Cambridge: Harvard University Press, 2010.
[4] H. Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life,* Palo Alto: Stanford Law Books, 2009.
[5] S. Shapiro, "Privacy Risk Analysis Based on System Control Structures: Adapting System-Theoretic Process Analysis for Privacy Engineering," *2016 IEEE Security and Privacy Workshops (SPW),* San Jose, CA, pp. 17-24, 2016.

17. Whether any aspects of the Cybersecurity Framework could be a model for this Privacy Framework, and what is the relationship between the two frameworks.

Having worked in both privacy and cybersecurity (including the Cybersecurity Framework), we believe the following aspects of the Cybersecurity Framework would be beneficial for the Privacy Framework:

- Employing multiple levels of abstraction, with a top level that non-experts can understand.
- Focusing on the "what" in the form of outcomes and activities, with pointers to resources to help determine the "how." This provides flexibility and enables mature organizations to map their current activities against the framework and less mature organizations to find useful resources for making progress. There is a growing body of work that would lend itself to inclusion as "informative references," some of which is mentioned in our other responses. Advancing this set of privacy informative references might also encourage the development of formal standards that incorporate it.
- A mechanism for tailoring the full menu of privacy content and mapping it to mission/business objectives, much like Profiles are used to get the most out of the Core in the Cybersecurity Framework. Linking activities and outcomes directly to mission/business objectives lends necessary context to strategic planning.
- Encouraging organizations to understand how they are situated within the larger environment. Much like the Cybersecurity Framework encourages organizations to understand their position in both the supply chain and critical infrastructure, the Privacy Framework could motivate organizations to understand the role(s) they play in the data ecosystem.

18. Please describe your preferred organizational construct for the Privacy Framework.

The possible options offered for consideration in the RFI obscure key distinctions between substantive, contextual, and context-free organizational schemes. FIPPs and the NIST privacy engineering objectives are examples of schemes which are grounded in concepts with specific substantive meaning. The information life cycle, use cases, and design patterns are examples of schemes that, while providing specific context, carry no substantive meaning of their own. In other words, these schemes orient thinking in a particular way, but they do not ground that thinking in specific semantics the way the FIPPs and the privacy engineering objectives do. The functions, categories, and subcategories of the Cybersecurity Framework constitute a scheme providing neither substance nor context. Rather, it is a purely structural scheme (a hierarchical taxonomy) with no implications for content whatsoever.

These different types of schemes represent a fundamental trade-off between conceptual leverage and flexibility. A substantive scheme, by utilizing a structure that itself carries meaning, would enable that meaning to be directly leveraged in the articulation and application of the Privacy Framework, but, by the same token, would constrain the concepts that could be expressed. A context-free scheme, on the other hand, would offer broad conceptual flexibility (its structure being the only constraint), but could not intrinsically advance the starting point of Framework development or application due to the lack of any "built-in" concepts. Contextual schemes, as one might expect, offer a middle ground that effectively splits the difference. Therefore, we

consider this question to be premature in the sense that it begs the more fundamental question of how much semantics to embed in the Framework's organizational structure.

All else being equal, it is probably more effective for purposes of both development and application to split the difference between substantive and context-free schemes, opting for a contextual scheme. This would enable the Framework to benefit from both some degree of conceptual leverage and some degree of flexibility, without completely sacrificing one of the other. However, we urge NIST to think in terms of a broad metamodel rather than a narrower orientation as represented by the examples of contextual schemes listed in the RFI. That meta-model should be agnostic with respect to substance but add value by establishing common organizational principles.

To achieve the desired attributes, particularly adaptability, ERM integrability, and compatibility, the Framework must be able to accommodate an increasingly heterogeneous substantive environment, including varying levels of organizational maturity, while adding value through contextual orientation. Relevant examples of this approach include the recently released OMG structured assurance case metamodel[6] and the ISO/IEC metamodel for software development methodologies[7] (which was the inspiration for a proposed metamodel for privacy engineering methods[8]). Based on trade-offs, environment, and precedent, we urge NIST to seriously consider this approach to structuring the Framework.

Note that a metamodel may be situated at different levels of abstraction. Therefore, depending on its ambitions for the Framework, NIST may want to consider a higher-level approach, such as the one depicted below.

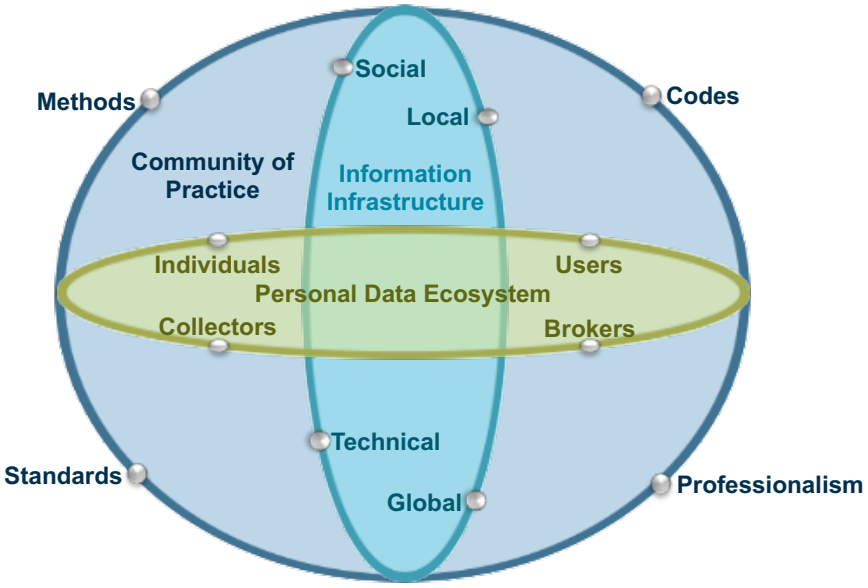

[6] Object Management Group (OMG), "Structured Assurance Case Metamodel (SACM), Version 2.0," Needham, MA: OMG, March 2018.
[7] International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) JTC 1/SC 27, "ISO/IEC 24744:2014, Software engineering—Metamodel for development methodologies," Geneva: ISO, 2014.
[8] Y.-S. Martin and J. M. del Álamo, "A Metamodel for Privacy Engineering Methods," *Proceedings of the 2017 International Workshop on Privacy Engineering,* San Jose, CA, 2017.

This privacy "biome" metamodel encompasses an ecosystem of privacy praxis in terms of three major dimensions: community of practice, personal data ecosystem, and information infrastructure. As such, it provides a kind of roadmap that NIST could employ to establish a Privacy Framework writ large, selectively developing more specific subsidiary metamodels for individual dimensions or even their constituent elements.

23. Whether some of these practices are inapplicable for particular sectors or environments.

We believe it is important to make no categorical presumptions of inapplicability. The sheer variety of specific contexts makes it impossible to produce anything resembling a definitive set of such exceptions. Moreover, such presumptions may create the potential for organizations to game the Framework, claiming use or adherence while exempting themselves from relevant practices. Claims of inapplicability should be made on a case-by-case basis and explicitly justified as part of Framework use.

Thank you again for the opportunity to contribute to the development of NIST's Privacy Framework. Please don't hesitate to contact us with any questions at sshapiro@mitre.org and jsnyder@mitre.org.

Sincerely,

Stuart Shapiro
Principal Cyber Security and Privacy Engineer

Julie Snyder
Lead Cyber Security and Privacy Engineer

The MITRE Corporation