From: rjc@enterprivacy.com <rjc@enterprivacy.com>
Sent: Thursday, October 24, 2019 1:24 PM
To: privacyframework <privacyframework@nist.gov>
Subject: NIST Privacy Framework: Preliminary Draft Comments

Please see attached PDF.

Jason

.....
R. Jason Cronk | Juris Doctor
Privacy Engineer | IAPP Fellow of Information Privacy
Enterprivacy | CIPT, CIPM, CIPP/US, PbD Ambassador
Privacy notices made simple: https://simpleprivacynotice.com
.....

October 24th, 2019

Katie MacFarland
National Institute of Standards and Technology,
100 Bureau Drive, Stop 2000
Gaithersburg, MD 20899
Submitted electronically to privacyframework@nist.gov

Comment on NIST Privacy Framework Preliminary Draft

Please find below my comments on the NIST Privacy Framework Preliminary Draft released September 6th, 2019. Overall, the Privacy Framework looks to provide a robust tool for organizations to build a privacy program and improve the privacy of consumers, employees and other individuals affected by organizational operations. The following comments are submitted to help improve the draft and future releases of the Privacy Framework to make it an even more valuable tool.

Clarification of Data and Data Actions as they relate to non-personal data

From my discussions with NIST, I understand that NIST is taking a broad view of the definition of 'data' and subsequently 'data actions' and 'data processing.' In particular, data includes non-personal data, unrelated to individuals or even groups of people. Indeed, the definitions in the Glossary make no distinction around personal versus non-personal data.

Data	A representation of information, including digital and non-digital formats.
Data Action (Adapted from NIST IR 8062 [5])	A system/product/service data life cycle operation, including, but not limited to collection, retention, logging, generation, transformation, use, disclosure, sharing, transmission, and disposal.

NIST correctly concludes that privacy risks can arise from actions performed on non-personal data. However, this reading will not be obvious to most casual users of the NIST Privacy Framework and most privacy professionals will instinctively interpret “data” as “personal data” given the historical connection between the processing of personal data and privacy. If NIST wishes organizations to interpret data and data actions broadly, that must be explicitly called out within the Privacy Framework.

Even beyond what most people would consider non-personal data, code is a form of data, which when executed affects a system’s state. Code can also create privacy risks. However, code is rarely thought of in terms of having data life cycle (collection, retention, etc.), thus where code creates risks absent the use of other data, it might not fall within the definition of data action, even though broadly defined. A better way to think about risks created by code, and possibly by non-personal data, is through interactions between systems and people (see next comment below).

Removal of 'Interaction' from the definition of Identify

NOTE: This comment points out deficiencies whereby the NIST Privacy Framework Preliminary Draft fails to

1.a.ii adequately strengthen individuals’ privacy protection

1.d. adequately define the relationship between privacy and cybersecurity risk

2.c. enable organizations to adapt to privacy risk arising from emerging technology such as the Internet of Things and artificial intelligence

In previous discussion drafts, NIST had defined the IDENTIFY-P (ID-P) function as "[d]evelop the organizational understanding to manage privacy risks for individuals arising from data processing **or their interaction with system, products or services."** As noted in my previous commentary submitted on the discussion drafts, while the IDENTIFY function included identification of risks related to interactions, the remaining functions in the Core dealt exclusively on managing privacy risks of data actions/processing. Rather than extend the other functions to include this important view of privacy risk, in the Preliminary Draft, NIST has chosen to strike the phrase referencing interaction from the definition of IDENITFY, which now reads: “[d]evelop the organizational understanding to manage privacy risks for individuals arising from data processing.”

As noted in my previous, unpublished commentary, interactions are an important factor in non-information privacy violations. There are many normative constructions of "privacy" and most include non-information privacy violations (**bolded**).

- Westin's four states of privacy: **solitude**, intimacy, anonymity and **reserve**;
- Prosser's privacy torts: **intrusion on seclusion**, public disclosure, false light and appropriation;
- Hartzog's three pillars of **autonomy**, **obscurity** and **trust**; and
- Solove's taxonomy which categorize privacy into Information Processing, Information Dissemination, **Collection (Surveillance and Interrogation) and Invasions (Intrusion and Decisional Interference).**

In the United States, many laws, regulation and constitutional interpretations deal with violations of non-information privacy interests, such as

- Right to privacy under the US Constitution (See *Griswold v Connecticut* ruling that state of Connecticut could not interfere with a family's private decision making regarding when to have children and their use of contraceptives)
- CAN-SPAM (which governs not the use of email addresses but rather the intrusion into individual's email inbox with unsolicited solicitations)
- Telemarketing Sales Rule which governs not the use of phone numbers but rather the intrusion into an individual's quiet enjoyment not to be disturbed (governing hours of calling, number of calls, and respect for opt-out decisions from being called). (See *Van Patten v Vertical Fitness Group*, LLC 847 F.3d 1037, 1041-43 (9th Cir. 2017) where the court ruled TCPA was established to protect right to be free from unsolicited telemarketing phone calls or text messages that "invade the privacy and disturb the solitude of their recipients.")
- The Wiretap Act prohibits wiretapping of phone conversations not the use of information gleaned from those conversations because wiretapping is an invasive act of surveillance in violation of social norms around listening in on people's conversations (see the etymology of *eavesdropping* for historical context).

Information/data privacy is certainly an important and necessary consideration in our information economy, but the future of privacy invasions rests in non-information privacy. And, while it is true that data can be used to facilitate non-information privacy invasions (just as a car can be used to get-away from a bank robbery but that doesn't make it a vehicular crime), governing data is a poor proxy for governing these types of invasions. Using Dan Solove's taxonomy, I'll provide four different scenarios

1. **Surveillance** – Ubiquitous cameras and microphones are creating a surveillance society. But it's not about the information that these collection devices collect, or the use of that information, it's the surveillance itself that changes people's behavior. This is what Ryan Calo calls the "perception of unwanted observation." It's why dummy

surveillance cameras work or Bentham's Panopticon. The potential of being observed changes behavior. In some cases, this has a positive social benefit (deterring crime) but in others cases it can have a negative social effect (deterring political organization). If all you have to do is give people the impression they're being observed to alter their behavior (to your benefit) how do you manage that with data processing controls?

2. **Interrogation** – Potential employers abusing the hiring process by asking irrelevant questions to a candidate will put the candidate at a disadvantage. It's not whether the employer uses the information (say about a candidate's personal relationship) to their advantage; the act of asking alters the perception and mental state of the candidate. Given the power dynamic, the candidate may question their ability to resist answering, lest they be overlooked. The privacy risk here isn't regarding a problematic data action, it's a problematic interaction.[1]

3. **Intrusion** – As more ubiquitous computing manifest itself through physical interactions with the environment, cyber physical intrusions into individuals' solitude and personal spaces will become more frequents. One could imagine a drone invading people's houses or hovering near people. What about robot sales solicitors that knock on every house on a city block trying to sell the occupants ice cream on a hot day? Even an internet connected baby monitor could facilitate the cyber physical intrusion of an Internet troll yelling obscenities at a baby (the physicality being the auditory emanations). This isn't problematic data action but it is an intrusion into the tranquility and sanctity of one's home and family. Consistent with this, I previously proposed in unpublished comments, a remapping of the privacy-cybersecurity risk relationship as displayed in the figure below (further commentary on definition of privacy risks

[1] One could stretch and say that the question "What's your current personal relationship status?" is a form of data and the problematic data action is use of that question/data in an interview. However, this is not an intuitive interpretation and not how most professionals would apply the Privacy Framework.

associated with data processing in the comment below).

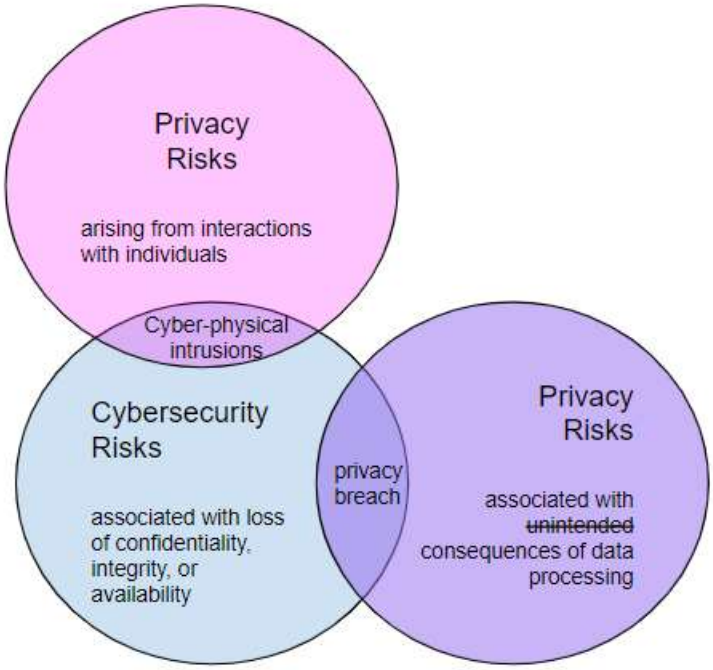

4. **Decisional Interference** – This is perhaps the most important privacy risks for the foreseeable future. While, analogous to the car use to facilitate a bank robbery, data can be used to facilitate manipulation of people's decisions (see EPDS Opinion on online manipulation and personal data), it needn't be the case (see Deceived by Design from the Norwegian Data Protection Authority). UI and UX design can take advantage of known psychological techniques to manipulate people to their disadvantage. This affects people's autonomy and erodes trust (see Hartzog's Privacy's Blueprint for more). It can be based on known techniques ('nudging') and cognitive biases broadly held in the population, and not on individual psychological weaknesses. It would be a stretch to consider a UI design used to manipulate people's decisions as a problematic data action in the way NIST has defined such.

While it seems clear than NIST doesn't want to tackle this important aspect of privacy at this time (given the striking of the interactions clause from the definition of IDENTIFY), I would suggest two alternatives to ensure that the Privacy Framework does not ignore non-information privacy risks. The first is to include an acknowledgement of the exclusion and desire to correct the problem in future versions in Appendix F, the Roadmap. The second would be to discuss it in terms of potential extensibility to the framework. Multiple times in the preliminary draft, NIST suggests the opportunity to add additional functions, categories and sub-categories but doesn't give insight into how or why the framework might be lacking in comprehensiveness. This would make a perfect opportunity

to include discussion of where an organization might want to extend the framework to address non-information privacy related issues that arise from their interactions with individuals.

Privacy Risks as 'unintended' consequences

NOTE: This comment points out deficiencies whereby the NIST Privacy Framework Preliminary Draft fails to

1.a.ii adequately strengthen individuals' privacy protection

1.d. adequately define the relationship between privacy and cybersecurity risk

In the previous discussion draft, the Venn diagram illustrating the difference between cyber-security risk and privacy risks, the privacy risk world was shown as representing "a byproduct of unauthorized data processing." In the new diagram, the privacy risk world is shown to be "associated with the unintended consequences of data processing." Unfortunately, this fails to capture that some privacy risks are associated with intended consequences of data processing. Some of these intended consequences may be socially acceptable (and within the risk tolerance of the organization or society) while some may not be.

If I'm a technology company designing software for a government agency to monitor citizens' social media posts, the numerous privacy violations (surveillance, aggregation of citizen data, identification of individuals, and others) and the subsequent secondary consequences (such as reduced discussion of politically sensitive topics and imprisonment of rabble-rousers) are clearly intended results of my data processing activities and clearly represent privacy risks to the citizens. But they are well within my organization's acceptable risk tolerance. I, therefore, suggest striking the word 'unintended.'

However, this also points to a more fundamental philosophical challenge within the framework. There is a notable lack of consideration for the legitimacy of data processing activities. Throughout the framework, there is an unstated presumption that the data processing will be performed and that we should constrain it with controls, but never question its validity. In my previous unpublished commentary, I suggested the inclusion of data minimization, noting it was only mentioned in passing regarding auditing and log files. I do note that in discussion draft 3, there was a category of Data Minimization, but none of the sub-categories really supported that title and it has been subsequently renamed Disassociated Processing.

Data minimization is such a core privacy practice; its absence from the framework is glaring. The Core does include CT.DP-P6 "Data processing is limited to that which is relevant and necessary for a system/product/service to meet mission/business objectives" which to the untrained eye may appear to be espousing data minimization but it doesn't. Article 5.1(c) of the European Union's General Data Protection Regulation (GDPR) essentially defines data minimization as "adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed." While this may seem similar to CT.DP-P6, it isn't. Under GDPR the purposes for processing must be lawful (see Article 6). Absent consent, contract, compliance, vital interest or in the public interest, processing to "meet mission/business objectives" must be legitimate. CT.DP-P6 includes no such qualification on mission/business objectives. The framework must support organizations being able to reconcile whether their objectives and the means by which they obtain those objectives are fundamentally flawed from a privacy risk perspective (i.e. lack legitimacy).

This comment is not unknown to NIST. Previously MITRE Corporation submitted a similar comment on the Draft Privacy Engineering Objectives and Risk Model.

> *The fact that a privacy attack is malicious does not mean it is automatically transformed into a security issue rather than a privacy issue, though it could be both. Consider, for example, dating website OkCupid's experimentation on its users, including putting the "wrong" people together to see if they would connect. This was arguably a malicious privacy violation, in that it was intentional. A proper human subjects research review process might have resulted in a mitigating protocol or even a decision not to proceed, but security principles were not relevant to this particular issue. We recommend the following edits as a potential way to address some of these comments, as shown in yellow highlight and strikethrough:*
>
> *Scope: The privacy engineering objectives and risk model are primarily focused on mitigating risks arising from ~~unanticipated consequences of normal system behavior~~ the behaviors of systems and their users.*[2]

There point was the same as mine; Privacy Risks (which impact individuals) may be perfectly intended and anticipated by the organization, but within their risk tolerance. I recommend the word "unintended" be stricken from Venn diagram in figure 2.

[2] MITRE Comments on the Draft Privacy Engineering Objectives and Risk Model, MITRE Corporation, Oct 15, 2014.

Consequentialist view of privacy risk

NOTE: This comment points out deficiencies whereby the NIST Privacy Framework Preliminary Draft fails to

1.a.ii adequately strengthen individuals' privacy protection

The preliminary draft takes a problematic view of privacy risks. While laudable in its attempt to deal with external risks an organization might impose on individuals, rather than a strict enterprise risk view, the view of privacy risks based on the impact it imposes on individuals causes several problems.

Not all organizational risks flow from individual impact

Figure 3 and the supporting text suggest that organizational risk from privacy stems from individual impacts (embarrassment, discrimination, economic loss). However, many organizational risks (regulatory fines, legal action, and eroded trust in the market) stem from violations of social norms and laws not injury or harm to the individuals. See, for example, the Illinois Supreme Court ruling (January 2019) in the *Rosenbach v. Six Flags* on the Illinois Biometric Information Privacy Act (BIPA):

> *When a private entity fails to adhere to the statutory procedures, as defendants are alleged to have done here, the right of the individual to maintain their biometric privacy vanishes into thin air. The precise harm the Illinois legislature sought to prevent is then realized. This is no mere "technicality." The injury is real and significant.*

This was reaffirmed by the US District Court of appeals on August 8th, 2019 in *Patel v. Facebook* where they ruled that violation of the substantive right to privacy in biometric information constituted injury-in-fact without the need for subsequent consequences to the individual plaintiffs. Even where the laws haven't codified privacy rights, social norms around behavior can create organizational trust issues prior to legal sanctioning of those behaviors.

Creates disincentives to transparency by organizations

By focusing privacy risks on consequentialist impacts to individuals, the NIST privacy risk analysis creates perverse incentives for organizations. The proposed risk model includes (starting at line 824)

> *Likelihood is defined as a contextual analysis that a data action is likely to create a problem for a representative set of individuals. Context can include organizational factors (e.g., the public perception about participating organizations with respect to privacy), system factors (e.g., the nature and history of individuals' interactions with the*

*system, **visibility of data processing** to individuals and third parties), or individual factors (e.g., individuals' demographics, privacy interests or perceptions, data sensitivity).*

As noted in the preceding section, "public perception about participating organizations" isn't an impact on individuals but rather an impact on the organization. One could argue that it does impact individuals because it alters behavior (avoiding the organization's products and services) but this might be a positive impact on the individual (by avoiding a bad product), not a "problem for individuals" or adverse consequence. Regardless, most organizations are not going to take this nuanced approach. They will focus on traditional notions of adverse consequences to the individuals: physical and mental injury.

The biggest problem with focusing on individual and subsequent organizational impact is it creates disincentives for transparency (hence the bolded phrase in the quoted text). By being honest and transparent about what the organization does, it could negatively impact the individual's psychological state, mental well-being or behavior. We've already seen this with Facebook where they tried to hide the collection of information in their mobile app by suppressing the permissions request (in order to avoid the public backlash), Uber's paying off a security researcher as a bug bounty to avoid disclosure of a security breach and with other companies as well. If my privacy risk model is purely consequentialist based then CM.AM-P1 seems at odds with reducing that risk, because communicating what an organization is doing increases the likelihood of mental state changes in individuals and increases the likelihood of regulatory action.

Doesn't adequately capture privacy risks

One of the key focuses of the privacy framework is the data ecosystem and third party organizations. Having a focus on individual impact ignores an important calculation in privacy risks; the likelihood of a "problematic data action" arising in the first place. Now, for the purposes of internal analysis, one could presume the inevitability of certain data processing activities by your own organization (which I think from a privacy design perspective is ill-advised but …). But at least as far as analysis of third parties within the ecosystem, incorporating the likelihood that a third party will engage in a substantive privacy violation (say secondary use of information) is imperative. Organizations may wish to forgo engagement with third parties likely to operate in ways that run counter to their privacy values. But the likelihood of such actions must be part of the privacy risk calculation. As previously mentioned this may play into anti-transparency incentives. For instance, an email provider may proactively supply information to certain government agencies, but for the vast majority of customers it won't create any adverse actions against

them. However, if it was learned that such information were being disclosed, it might cause psychological trauma to a large population of the customers, who thought their email was safe from the prying eyes of government agents. Therefore, the organization decides to obfuscate such disclosure with ambiguous language in their privacy notice. Under the consequentialist privacy risk analysis, they've implemented a perhaps uncodified sub-category of an anti-transparency outcome. In other words, they've reduced risks both to customers (making it less likely that customers will be freaked out and upset) and organizational risks (making it less likely that the customers will go to a different service). Instead, what they should have analyzed was whether providing such information to the agency would violate a social norm and expectation of their customer base. In other words, the organization should reduce the actual likelihood of a privacy violation and the severity of that violation, not the likelihood and magnitude of the consequences.

Viewed from the larger lens of privacy risk including risks of privacy violations, NIST's consequentialist approach seems unnecessarily narrow.

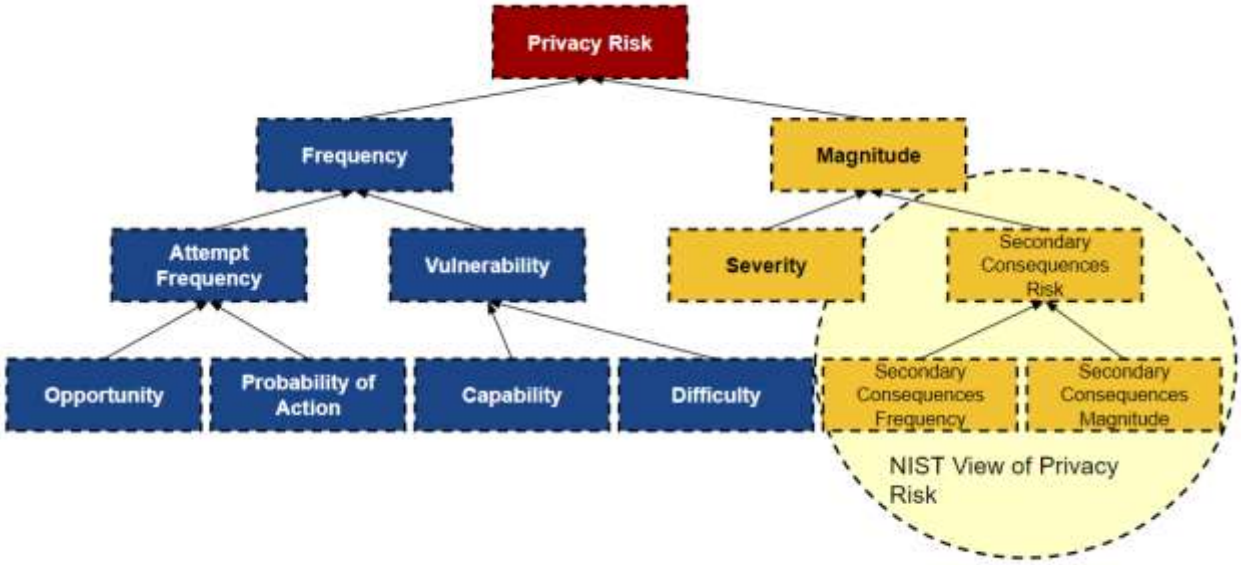

Similarly, MITRE Corporation has provided NIST with feedback on the difficulty of discussing harms in a community not steeped in the literature on risk management or privacy harms. While the concept of violations/harms maybe harder to see in an Information Privacy context, when considering non-information privacy violations, such as surveillance, the notion seems much more obvious or intuitive. Consider a covert surveillance camera placed in someone's house. This is a clear violation of social norms around privacy and the sanctity of one's home. We don't need to use the information, or further disseminate it; the act of surveilling is the violation. In this respect, we're not

concerned about the secondary consequences (the impacts such as embarrassment). Even absent this, we can assess some measure of severity of the violation. Is it in a common area, such as a living room? How about a bedroom? How about a bathroom? The latter seems a more severe invasion of privacy, because of our social mores around privacy in these types of internal rooms in our already private residence. Some will argue that the concern here is the risk of discovery (and thus that impact of embarrassment). Yes, those risks of secondary consequences exist. The point is, that even if negligible, the severity of the initial violation MUST be considered when assessing privacy risks.

There are three solutions to the problem presented above (in order of how best to address it).

1. NIST could redefine the privacy risk analysis to incorporate the likelihood that a data action is problematic because it violates laws or social norms and the severity of that violation. From these violations of social norms could stem additional secondary risks (with their own likelihoods and impacts to individuals, the organization and society).
2. NIST could expand on its notions of 'privacy values. ' I think privacy values is an imperative but under-explored concept in the Privacy Framework and bears much more discussion. This is especially true given the consequentialist risk model. Unfortunately, most organizations are likely to overlook the privacy values component of the analysis and simply use the Core as a checklist, as NIST cautions against. Right now, privacy values pulls a lot of weight despites relegation to a few mentions throughout the document. Expounding on privacy values, suggesting sources identifying values, and discussing the risks of the organization (or third parties) engaging in data actions in contravention of the identified privacy values could be used as a precursor to the proposed privacy risk model to include mitigation of privacy violations absent subsequent harm.

   The problem with the privacy values approach is that given an organization defines its own values, which may contravene social norms and laws, they expose individuals to unanalyzed privacy risks, when they should be acknowledging the risks and placing them within acceptable risk tolerance levels for their particular organization.
3. Some have suggested the definition of impact or harms incorporate violations of social norms and laws, as the US courts seem to do. I would suggest this is problematic, because first it doesn't square with the dictionary definition of harm (which relates to physical or mental injury) and it doesn't match most people's perceptions around "harm" and thus will lead to confusion in implementation.

Sincerely,

R. Jason Cronk
Author, *Strategic Privacy by Design*
Principal Privacy Engineer
Enterprivacy