Federal Register Notice 87 FR 35250, https://www.federalregister.gov/documents/2022/06/09/2022-12432/request-for-information-on-advancing-privacy-enhancing-technologies, June 9, 2022

Request for Information (RFI) on Advancing Privacy Enhancing Technologies

Immuta

Immuta's Response to OSTP's Request for Information on Advancing Privacy-Enhancing Technologies.

Immuta is the leading universal cloud data security platform, providing one platform to automate access control for any data, on any cloud service, across all compute infrastructure. Immuta also offers advanced privacy-enhancing technologies (PETs) to support data re-use and data sharing use cases.

At Immuta, we welcome the effort to advance privacy-enhancing technologies, which should be informed by a rigorous analysis of what PETs actually achieve in terms of privacy and security. We suggest that a risk-based approach to PET adoption is key to put data controllers or processors in a position to choose the right PET or set of PETs for each data re-use or data sharing use case with which they are confronted. In fact, more often than not, PETs should be considered as complementary rather than mutually exclusive. Through this response, we would like to share a summary of our PET evaluation framework[1] as one such risk-based approach. We have found such a framework helps stakeholders identify risks related to the implementation of a variety of PETs in the context of a specific use case.

Our PET evaluation framework relies upon threefold considerations supporting the claim that not all PETs are equal in all circumstances.

1. PETs don't offer the same formal (mathematical) guarantees. Generally speaking, soft PETs don't offer any formal guarantee, while hard PETs do. Still, these guarantees are not necessarily related to the same type of privacy or security risk, even if many of these PETs aim to achieve the same privacy or security goal, i.e. confidentiality.[2] By way of

[1] Presented at the Future of Privacy Forum masterclass on de-identification on 25 May 2022 at the Computer, Privacy and Data Protection Conference 2022, available at https://www.youtube.com/watch?v=r1J93rTeeaY.

[2] Privacy and Security requirements are usually broken down into subgoals in order to help with risk assessment and/or compliance. Security goals traditionally include confidentiality, integrity and availability, while privacy goals are usually considered to be more numerous and thus cover, on top of the

example consider two hard PETs, one PET could guarantee that every participant in a data set is in a position to deny the content of its record (this is the formal guarantee achieved by local differential privacy), while another could guarantee that within a data source there is at least a k number of identical records, so that no individual record is actually unique (this is the guarantee achieved by k-anonymization). It's important to remember that formal guarantees are intimately related to threat models. Each formal guarantee holds for a predetermined threat model.[3]

2. PETs may or may not assume a trusted data processor, i.e., an entity in direct contact with the data, for example because it is transforming the data or because it is running computation over the data. The hard/soft distinction is therefore a useful starting point to determine what a specific PET achieves in the context of a particular use case. Soft PETs often assume a trusted data processor. Assurances are thus mainly derived from organizational controls (including legal obligations, monitoring and auditing processes). On the other hand, hard PETs are considered robust to misbehaving, malicious, and negligent data processors. Assurances are therefore mainly derived from mathematics. Often, no second or third party protocol participant is deemed to be in a position to violate relevant privacy or security interests, such as confidentiality. As a rule of thumb, soft PETs include the following PETs: Synthetic Data, Access Controls, Encryption (internal, with key management challenges)[4], Monitoring, and Auditing. Hard PETs include: Secure Multi-party Computation, Zero Knowledge Proofs, Partially Homomorphic Encryption, Fully Homomorphic Encryption, Differential Privacy, Shuffle Privacy, Local Differential Privacy, and Encryption (external, with no key management challenge).

---

security triad, data minimisation, purpose limitation or unlinkability, transparency, intervenability, fairness and accountability.

[3] See e.g. Marsh, Catherine, et al. 'The Case for Samples of Anonymized Records from the 1991 Census'. Journal of the Royal Statistical Society. Series A (Statistics in Society), vol. 154, no. 2, 1991, p. 305. DOI.org (Crossref), doi:10.2307/2983043.; Dankar, Fida Kamal, and Khaled El Emam. 'A Method for Evaluating Marketer Re-Identification Risk'. Proceedings of the 1st International Workshop on Data Semantics - DataSem '10, ACM Press, 2010, p. 1, doi:10.1145/1754239.1754271.

[4] Operational risks of misusing the key are secondary to mathematical guarantees if an organization can decrypt. Therefore, internal use of encryption requires operational assurances.

3. PETs may be utilized in combination to enhance their properties. Let's take the example of Federated Learning, which attempts to train a shared global model from local model updates. It is a soft PET but there are ways to harden it. While updating a global model running upon an organization's server from local models running on edge devices could seem a good idea because only aggregate data is exchanged between the organization's server and the edge devices, there is no guarantee that the push to the organization's server does not reveal anything too private or personal. This is because aggregates can be rich, and can memorize or signal in completely unexpected ways the participation of an outlying individual for example. Federated learning can thus be hardened by applying a hard PET such as local differential privacy.[5]

Our PET evaluation framework produces a risk profile for each use case in order to inform the choice of the right PET or set of PETs for the use case at hand. Two to three types of stakeholders are considered to perform the risk assessment: a data processor, who performs tasks upon the data usually on behalf of the entity initiating the processing or its collaborator; a collaborator, who collaborates with the entity initiating the processing because it is interested in the outputs of the processing, i.e., the query results; and a model user, who is involved when the processing output is a model and the collaborator discloses the model to a third party that is different from the entity initiating the processing or the data processor. Our analysis makes use of the attack model used in differential privacy[6] as a general framework for analyzing data

[5] Differential privacy, which is a family of mathematical techniques that formally limit the amount of private information that can be inferred about each data subject. There are two main flavors of differential privacy, offering slightly different privacy guarantees: "global," which offers data subjects deniability of participation, and "local," which offers deniability of record content. Despite being slightly different, both operate by introducing randomization into computations on data to prevent an attacker from reasoning about its subjects with certainty. Ultimately, these techniques afford data subjects deniability while still allowing analysts to learn from the data.

[6] See e.g., Dwork, Cynthia, and Aaron Roth. "The Algorithmic Foundations of Differential Privacy." Foundations and Trends® in Theoretical Computer Science 9, no. 3–4 (2013): 211–407. https://doi.org/10.1561/0400000042; Stalla-Bourdillon, S., & Rossi, A. (2021). Aggregation, Synthesisation and Anonymisation: A Call for a Risk-based Assessment of Anonymisation Approaches. In D. Hallinan, R. Leenes & P. De Hert (Eds.). *Data Protection and Privacy: Data Protection and Artificial Intelligence* (Computers, Privacy and Data Protection, pp. 111–144). Oxford: Hart Publishing. Retrieved July 5, 2022, from http://dx.doi.org/10.5040/9781509941780.ch-005.

privacy controls and expands it to allow for additional assumptions on the capabilities and resources of the model user. This allows our framework to be inclusive of other PET threat models.

Each PET is then assessed through nine dimensions, as defined below:

- **Formal guarantee:** guarantee that can be mathematically established, e.g. deniability of the attribute values included in a record or deniability of participation to a data set with a set number of records.
- **Threat model:** model through which relevant threats, or the absence of appropriate controls, can be identified, and controls prioritized. It usually relies upon the conceptualization of a situationally-relevant attacker (e.g., a journalist, a marketer, a prosecutor, or an untrusted analyst) with situationally-relevant characteristics, such as access to publicly available information, access to the protected data, expert skills, and/or prior knowledge.
- **Loss of fidelity:** loss of accuracy (in comparison with the original version of the data) due to the transformation of the data with a view to mitigate risks, e.g. through the injection of noise.
- **Processor inference:** the likelihood a processor could uncover the identity of an individual, the values of (sensitive) attributes attached to an individual record, the participation of an individual to a data set, or the relationship or link between one or more individual records.
- **Processor sabotage:** deliberate subversion of the processing by the processor, e.g. through collusion with another participant.
- **Collaborator inference:** the likelihood the collaborator could uncover the identity of an individual, the values of (sensitive) attributes attached to an individual record, the participation of an individual to a data set, or the relationship or link between one or more individual records.
- **Collaborator sabotage:** deliberate subversion of the processing by the collaborator, e.g. through collusion with another participant.
- **Model user inference:** the likelihood a model user could uncover the identity of an individual, the values of (sensitive) attributes attached to an individual record, the

participation of an individual to a data set, or the relationship or link between one or more individual records.

- **Practical barriers:** practical barriers to the implementation of the PET, e.g. barriers related to computation speed or set-up complexity.

Formal Guarantee	Threat Model	Loss of Fidelity	Processor Inference	Processor Sabotage	Collaborator Inference	Collaborator Sabotage	Model User Inference	Practical Barriers

Table 1. The 9 dimensions of PETs

This framework thus shows that each PET usually has a narrow privacy focus, most of the time confidentiality, and often PETs must be deployed in thoughtful combinations in order to mitigate a variety of data processing risks, including processor inference and sabotage, collaborator inference and sabotage and model user inference.