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

Kulshrestha, Anunay; Mayer, Jonathan; and Scheffler, Sarah

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views or opinions of the U.S. Government. We bear no responsibility for the accuracy, legality, or content of the responses and external links included in this document.

PRINCETON UNIVERSITY

CENTER FOR INFORMATION TECHNOLOGY POLICY
PRINCETON UNIVERSITY

July 8, 2022

Office of Science and Technology Policy
Executive Office of the President
Eisenhower Executive Office Building
1650 Pennsylvania Avenue
Washington, DC 20504

Comment of Princeton University Researchers on Advancing Privacy-Enhancing Technologies

Thank you for the opportunity to inform the forthcoming national strategy on privacy-preserving data sharing and analytics. We are academic researchers at Princeton University who study information privacy and security using interdisciplinary methods. Our scholarship has examined data practices that affect individual privacy, the contours and efficacy of privacy law, and the design and implementation of new privacy-enhancing technologies.

In this comment, we briefly respond to two prompts in the request for information, drawing on our recent research experiences.

7. Risks related to PETs adoption: Identification of risks or negative consequences resulting from PETs adoption as well as policy, governance, and technical measures that could mitigate those risks.

Privacy-enhancing technologies can provide unprecedented capabilities, as the RFI rightly notes. But, in some instances, those capabilities can risk causing serious harms—calling into question whether the technology merits deployment.

In a recent publication, we constructed a novel privacy-preserving protocol for identifying harmful imagery in end-to-end encrypted services.[1] The conclusion of our

[1] Anunay Kulshrestha & Jonathan Mayer, *Identifying Harmful Media in End-to-End Encrypted Communication: Efficient Private Membership Computation*, Usenix Security Symposium (2021).

work was that, while the technology was viable, it also created a danger of misuse for surveillance and censorship by foreign governments.[2] Shortly before we were scheduled to present the publication, Apple announced a very similar privacy-preserving protocol for identifying harmful imagery stored with its iCloud Photos service.[3] A global backlash followed—including from experts in PETs—because of the possibility that the system could backfire and *undermine* rather than *enhance* privacy.[4] Concerns also emphasized possible threats to free speech and risks for vulnerable users. In response, Apple decided to pause its deployment of the PET.

We take two lessons from this experience. First, where a PET creates a capability that did not exist before, it is important to carefully examine the societal context for and consequences of that capability. Risks of misuse, exploitation, inaccuracy, or inequity may outweigh the upsides of the PET. These risks depend not only on the current implementation of the system, but also on the future motivations and capabilities of persons, organizations, and governments that could influence the system going forward. Communicating with the public about these risks may be extraordinarily difficult, and quantifying these risks may be effectively impossible.

Second, PETs can be more difficult to supervise and audit than conventional system designs, because they (intentionally) restrict the information that parties can obtain. Unfortunately, PETs research usually starts and ends with providing a new capability in a privacy-preserving way. There is an urgent need for complementary research on how to achieve transparency and accountability for PETs, while maintaining the strong privacy guarantees that make them so promising. In our own upcoming research, for example, we are developing new designs for public transparency in privacy-preserving hash matching systems.[5]

[2] Jonathan Mayer & Anunay Kulshrestha, *We Built a System Like Apple's to Flag Child Sexual Abuse Material—and Concluded the Tech Was Dangerous*, Wash. Post (Aug. 19, 2021).

[3] Abhishek Bhowmick et al., *The Apple PSI System* (July 29, 2021).

[4] *See* Hal Abelson et al., *Bugs in Our Pockets: The Risks of Client-Side Scanning* (Oct. 15, 2021); Sarah Scheffler & Jonathan Mayer, *Content Moderation in End-to-End Encryption* (in submission).

[5] Sarah Scheffler et al., *Public Verification for Private Hash Matching: Challenges, Policy Responses, and Protocols* (in submission).

3. Specific sectors, applications, or types of analysis that would particularly benefit from the adoption of PETs: Information about sectors, applications, or types of analysis that have high potential for the adoption of PETs.

The federal government has an opportunity to lead by example. By adopting PETs, the government can offer new forms of public transparency and more rigorous statistics than were previously possible. The Census Bureau's implementation of differential privacy stands out as a particularly remarkable example of what the federal government can achieve by adopting PETs and how government use of PETs can advance the entire field.

Transparency about the data that the government holds about individuals, and the privacy implications of that data, is another promising application of PETs. In recent research, we examined how secure multiparty computation could provide a viable path forward for the decade-long challenge of estimating how Foreign Intelligence Surveillance Act collection affects ordinary Americans.[6] While our work focuses on just one type of data (intercepted communications) and just one type of government component (elements of the Intelligence Community), the same concept could bring transparency to personal data holdings across the federal government.

With careful design, PETs could enable the best of both worlds: maintaining the necessary confidentiality and accuracy of important government systems, while also significantly improving their transparency and accountability.

* * *

Thank you again for the opportunity to provide input to the federal government's strategic plan for information integrity research and development. We would be glad to provide additional detail or discussion as would be helpful to OSTP.

[6] Anunay Kulshrestha & Jonathan Mayer, *Estimating Incidental Collection in Foreign Intelligence Surveillance: Large-Scale Multiparty Private Set Intersection with Union and Sum*, Usenix Security (2022).

Sincerely,[7]

Anunay Kulshrestha
Graduate Researcher, Center for Information Technology Policy, Princeton University

Jonathan Mayer
Assistant Professor of Computer Science and Public Affairs, Princeton University

Sarah Scheffler
Research Fellow, Center for Information Technology Policy, Princeton University

[7] We offer this comment as individual academic researchers.