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

Stealth Software Technologies Inc

Steve Lu
CEO

Stealth Software Technologies, Inc.
12100 Wilshire Blvd. 8th Floor
Los Angeles, CA 90025

July 8, 2022

Submitted via email to PETS-RFI@nitrd.gov

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

Re: Notice of request for information on Advancing Privacy-Enhancing Technologies

Dear Sir/Madam:

Stealth Software Technologies, Inc. ("Stealth") appreciates the opportunity to comment on the Office of Science and Technology Policy's request for information on Advancing Privacy-Enhancing Technologies (the "OSTP RFI"). Stealth is a small business co-founded by Dr. Rafail Ostrovsky, a Distinguished Professor of Computer Science and Mathematics at UCLA. Our team consists of experts in cryptography, computer science, and mathematics; experienced software engineers; and numerous consultants who are leading academic researchers. Stealth has performed on numerous Federal Government contracts over the past decade, for agencies and federal entities including DARPA, IARPA, ONR, and NIST.

Our comments below address several of the specific topics presented in the OSTP RFI, on a topic-by-topic basis.

Topic 1: *Specific research opportunities to advance PETs*

As a small selection of the many potential Federal research initiatives that could substantially accelerate and improve both the development and adoption of privacy-enhancing technologies (PETs), we identify some particularly productive opportunities for research programs and strategies to develop such programs:

- A primary barrier to adoption of PETs—particularly those based on modern cryptographic tools such as secure multi-party computation or homomorphic

encryption—is that, at present, even the *use* of such technologies to their full privacy-enhancing potential frequently requires highly specialized expertise and extensive training. A key contribution to overcoming this barrier would be the development of a *programming language and software stack for expressing privacy-preservation requirements* designed for users who are *not experts in cryptography*; such a programming language would enable users to obtain optimally privacy-preserving and application-appropriate cryptographic protocols while shifting the burden of cryptographic expertise to the developers of the language.

Toward this end, in 2019 the Intelligence Advanced Research Projects Activity (IARPA) created a broad agency announcement program titled "HECTOR" aimed at developing such a software stack. Both the HECTOR program itself and its source-selection process required a partnership between programming-language experts, systems experts, and cryptography experts. Although the four selected performer teams were first-rate and the first year of execution resulted in substantial accomplishments, the program was unfortunately canceled. We believe that resuming such an interdisciplinary research and development effort would substantially advance the field of PETs—particularly with a view toward the critical aspects of *usability* and *accessibility*.

- A frequently encountered limitation of PET tools—particularly those that provide *exact* computations with strong privacy guarantees through cryptographic protocols—is that their memory requirements are currently too great for many common real-world applications. Such limitations arise particularly often in highly computationally intensive applications to artificial intelligence and machine learning on large data sets, and a targeted research effort to enhance the performance of cryptographic PETs, such as secure multi-party computation and fully homomorphic encryption, on large data sets is very much needed.

  A primary reason for this excessive memory consumption is that such PETs usually require computer programs to be converted into digital circuits for secure execution. Theoretical techniques enabling random access memory (RAM)-based secure computation, thereby avoiding the need for conversion to circuits, do exist; however, practical solutions for implementing these theoretical techniques and integrating them with existing data systems and infrastructure—such as large data-base management systems like Apache Spark, large training sets for artificial intelligence and machine learning applications, data analytics platforms, and data visualization formats—are missing. Again, research targeted toward such practical implementation of cryptographic theory are critical to advancing the usability and accessibility of PETs.

- The Defense Advanced Research Projects Agency (DARPA) has been instrumental in administering programs advancing fundamental and applied research in cryptography for the purpose of enhancing privacy. As evidence of its impact, leading PETs including work in secure multi-party computation such as Prio, Poplar, and Prio+; work in fully homomorphic encryption including PALISADE, and works in differential privacy that intersected with the 2020 decennial United States Census have all been funded in part by

various DARPA programs. Although research advancing privacy and confidentiality fits naturally within the Department of Defense in the interest of national security, it is clear that privacy considerations present common to diverse agencies across the Federal Government. Programs for PET research, tailored toward the activities and goals of specific Government entities, should be more widespread and funded by agencies beyond the Department of Defense. To this end, actively recruiting program managers with expertise in privacy and cryptography into research agencies in other departments is crucial to ensure that PETs advance further.

Topic 2: *Specific technical aspects or limitations of PETs*

Decision tools for PET selection

A common difficulty voiced by interested potential adopters of PETs is that they do not know where or how to begin in their efforts to select optimal technologies for their particular applications. Due to the concurrent emergence of various cryptographic, statistical, and hardware-based approaches to PETs in recent years, there has been significant confusion regarding which PET techniques to use for which purposes. This confusion is well founded—indeed, the various techniques and tools comprising the wider category of PETs are meaningfully distinct, and PETs are not fungible; depending on the specific application scenario or requirements of a given use case, some PETs may not even be meaningfully applicable.

This difficulty is of course not unique to PETs—all systems are to some extent sensitive to subtle differences in requirements—but many of the underlying issues have largely been adequately resolved in other technological areas, in which principled, requirements-based solution selection can be accomplished using many basic tools that are currently lacking in the field of PETs. These include accessible, well-designed, and widely available decision aids including decision trees and flowcharts, guidelines and best practices, pillars of success, knowledge bases and knowledge diagrams, and general systematizations of knowledge; such tools can be compounded with prior successful approaches to applications through the use of white papers, case studies, and roadmaps. The field of PETs would be advanced significantly through the development of such decision aids for the adopting public.

Limitations of particular PETs

Each particular PET approach has drawbacks, and understanding these drawbacks is critical to inform the optimal solution selection addressed above. Differential privacy, for example, relies inherently on the introduction of noise to data in order to ensure privacy, resulting in some likely amount of distortion of computational output. Accordingly, differential privacy is an inappropriate technology for applications in which an *exact* computation is needed, and users seeking privacy in such applications should turn to secure computation frameworks such as secure multi-party computation or homomorphic encryption. These latter approaches guarantee that inputs to a computation remain hidden, while the exact result—and only the result—is revealed to specified parties; it is frequently not appreciated by general audiences that such

cryptographic techniques do not involve "fuzzing" of data to guarantee privacy. Differential privacy is similarly ill-equipped to handle database queries asking whether there exists a database entry satisfying a specific condition, and secure computation frameworks are again more appropriate. Inversely, for applications in which privacy is compromised by inferences that can be drawn from the exact results of an analysis, differential privacy is a more appropriate tool.

Finally, we note certain significant technical differences between secure multi-party computation and homomorphic encryption. So-called *fully* homomorphic encryption—which is required to handle *arbitrary* computations—necessarily requires a conversion to digital circuits for secure computation, whereas, as discussed in our response to Topic 1 above, there exist techniques for executing secure multi-party computation using random access memory. Additionally, homomorphic encryption requires key-management mechanisms and trust in another party holding the decryption key, whereas secure multi-party computation essentially accomplishes decryption "within the system" through a protocol, at the cost of requiring interactivity and in some cases trust assumptions of non-collusion.

Finally, with respect to hardware-based PETs such as secure enclaves, although the application of such technologies to software attestation has been a major success, maintaining privacy of computations inside enclaves has been an elusive goal, and such technologies have routinely been compromised by novel attacks.

Interaction of PETs and blockchain

Blockchain technology is another cryptographic paradigm that has significantly increased in visibility in recent years; it enables an immutable recording of data at a large scale, applicable not only to data held by a single entity but also across multiple organizations and data fields. A combination of the privacy afforded by PETs with the auditability afforded by blockchain technologies, when applied to Government records, could potentially enable substantial advances in preventing, identifying, and mitigating fraud and abuse—it should be stressed that although auditability and privacy may initially appear to be competing or even mutually contradictory goals, the appropriate use of PETs that enable *computation* on private data can enable these goals to be realized simultaneously.

Topic 4: ***Specific regulations or authorities that could be used, modified, or introduced to advance PETs***

In the near term, any new regulations or modifications to existing regulations mandating or otherwise controlling use of PETs should be carefully drafted to remain appropriately flexible and technology-neutral, in order to avoid prescribing the use of particular PETs before the technologies—and their optimal spheres of application—have been sufficiently standardized. Indeed, as addressed in our comments on Topic 6 below, given the relatively nascent and rapidly evolving nature of PETs as practically applicable technologies, rigid standardization itself may be premature at present. Moreover, until PETs become more accessible through widespread

adoption and awareness, strict regulatory mandates for their use in the private sector may be logistically or economically infeasible for many regulated entities. However, appropriately flexible initial measures short of highly specific regulatory mandates can further such adoption and awareness. In this direction, requirements in the Federal Acquisition Regulations requiring appropriate use of PETs by contractors in data management or in deliverables to the Government would be welcome, but—particularly prior to standardization—such requirements should be crafted to allow contractors to incorporate PETs feasibly and in a genuinely beneficial manner.

There may be significant opportunities to encourage PET adoption by issuing guidance clarifying the application of existing regulations to appropriate uses of PETs. Particularly in sectors covered by existing privacy-regulation regimes, disruption of established compliance frameworks and particularly uncertain risks of non-compliance can deter regulated entities from adopting novel technologies, which in many cases may actually provide highly effective technical measures for compliance with privacy regulations if applied appropriately.

Guidance from regulatory bodies can mitigate such deterrents, simultaneously advancing the adoption of PETs and facilitating effective compliance; it may be appropriate to issue guidance specifically devoted to addressing likely uses of PETs by regulated entities and the compliance status of such uses under applicable privacy and security rules.

In the context of financial regulation, the Federal Trade Commission (FTC) Privacy of Consumer Financial Information Rule[1] ("Privacy Rule") and Standards for Safeguarding Customer Information Rule[2] ("Safeguards Rule"), covering financial institutions under the Gramm–Leach–Bliley Act, are prime examples of regulations that could benefit from such clarifying guidance. As a particular example, the Privacy Rule generally prohibits covered entities from disclosing customer account numbers to nonaffiliated third parties for marketing purposes, but specifically provides that "[a]n account number . . . does not include a number or code in an encrypted form, as long as [the covered entity does] not provide the recipient with a means to decode the number or code."[3] The phrase "encrypted form" is not defined directly in the Privacy Rule; although the Safeguard's Rule's definition of "encryption" could arguably apply to, for example, the use of account numbers in a secret-sharing-based secure multi-party computation protocol,[4] there may be sufficient ambiguity in the language to deter a covered entity from considering such use—even if the entity were willing and able to undertake the combination of regulatory and technical analysis required to assess it. By contrast, the applicability of the Privacy Rule's exception to an analogous computation involving account numbers using homomorphic encryption, with substantially similar privacy and security guarantees, would be considerably less ambiguous due solely to terminology. We offer no opinion in this comment regarding how the Privacy Rule or Safeguards Rule *should* apply to any specific application of any particular PET. We simply note

[1] 16 C.F.R. § 313.1 (2021).
[2] *Id.* § 314.1.
[3] *Id.* §§ 313.12(a), (c)(1).
[4] *See id.* § 314.2(f) (defining "encryption" for purposes of the Safeguards Rule as "the transformation of data into a form that results in a low probability of assigning meaning without the use of a protective process or key, consistent with current cryptographic standards and accompanied by appropriate safeguards for cryptographic key material"); *id.* § 313.3 (containing no definition of "encrypted" for purposes of the Privacy Rule).

that guidance from the FTC clarifying the application of these rules to foreseeable applications of a diverse variety PETs would be welcome, particularly to the extent that it encourages beneficial uses of PETs.

In a potentially related vein, the President's July 2021 Executive Order on Promoting Competition in the American Economy[5] has encouraged the Consumer Financial Protection Bureau (CFPB) to consider rulemaking under section 1033 of the Dodd–Frank Wall Street Reform and Consumer Protection Act.[6] Section 1033 directly addresses portability and sharing of consumer financial transaction data rather than computation involving such data, and we offer no opinion on the scope of the CFPB's section 1033 rulemaking authority. However, to the extent that data-privacy considerations are implicated in any such rulemaking, they should be informed by an awareness of the potential applicability of PETs to sharing of and analytics and other computation on consumer financial data.

PETs may have beneficial applications to anti-money laundering (AML) and know-your-customer (KYC) activities, potentially allowing for both effective execution of AML and KYC policies and demonstration of compliance with such policies while preserving the privacy of sensitive data. Accordingly, rulemaking and enforcement bodies, notably the Financial Crimes Enforcement Network (FinCEN) and the Office of Foreign Assets Control (OFAC), should remain aware of the evolving state of the art in PETs and, if appropriate, issue guidance concerning such uses of PETs. As a specific example, PETs could potentially be relevant to FinCEN's Customer Due Diligence Rule, particularly with respect to the requirement that regulated institutions implement and maintain customer due diligence procedures for "[u]nderstanding the nature and purpose of customer relationships for the purpose of developing a customer risk profile" and ongoing risk monitoring.[7]

In the realm of health care regulation, PETs-specific guidance from the Department of Health and Human Services (HHS) similar to that proposed above from the FTC would be welcome, in the context of the HHS Privacy of Individually Identifiable Health Information Rule ("Privacy Rule") and Security Standards for the Protection of Electronic Protected Health Information Rule ("Security Rule")[8] under the Health Insurance Portability and Accountability Act (HIPAA). In particular, by facilitating joint computation involving protected health information (PHI) without requiring the direct sharing of such information, PETs can further the goals of the "minimum necessary standard" of the HIPAA Privacy Rule, which restricts the use and disclosure of PHI to that necessary to satisfy a particular purpose or carry out a given function.[9] However, uncertainty regarding the application of the HIPAA Privacy and Security Rules to particular applications of PETs may deter beneficial adoption of privacy-preserving analytical tools for treatment and research. As one particular example, the Privacy Rule allows a covered entity to use or disclose a "limited data set," defined to exclude certain direct identifiers of individuals, for

[5] 86 Fed. Reg. 36,987 (July 14, 2021).
[6] 12 U.S.C. § 5533 (2018).
[7] 31 C.F.R. § 1020.210(b)(2)(v) (2021).
[8] 45 C.F.R. §§ 164.102 *et seq.* (2021).
[9] *See id.* §§ 164.502(b), 164.514(d).

research, public health, or health care operations subject to a suitable data use agreement with the recipient of the limited data set.[10] Such direct identifiers may be particularly useful, perhaps even practically essential, to *link* data sets held by different covered entities; various PET tools and techniques can be used to achieve such linking without revealing any information about the underlying direct identifiers, significantly expanding the scope of available analytics for research and public health activities. However, it may be unclear whether such a use would be permissible under the regulations' definition of "use."[11]

Topic 5: ***Specific laws that could be used, modified, or introduced to advance PETs***

Although at present we have no specific proposals for new legislation that could be introduced, we support the existing bills introduced for the Promoting Digital Privacy Technologies Act,[12] the Digital Accountability and Transparency to Advance Privacy Act ("DATA Privacy Act"),[13] the Algorithmic Accountability Act,[14] and the National Secure Data Service Act[15] in furtherance of a national strategy for advancing PETs.

Regarding data sharing—or more accurately, in this case, data *processing*—among international entities, we specifically note guidance from the European Data Protection Board explicitly recognizing secure multi-party computation as an "effective supplementary measure" under the European Union's General Data Protection Regulation (GDPR) for data processing by independent processors located in distinct jurisdictions, as well as an effective technical privacy-by-design measure to ensure that data are not subject to unauthorized access.[16]

Topic 6: ***Specific additional mechanisms that could be used, modified, or introduced to advance PETs***

The National Institute of Standards and Technology (NIST) should play a central role in advancing the development and adoption of PETs in the United States (and, indeed, globally), as recognized in some of the proposed legislation discussed in our comments on Topic 5. As noted in our response to Topic 4, however, NIST should be careful not to set standards recommending

---

[10] *See id.* § 514(e).
[11] *See id.* § 160.103 ("*Use* means, with respect to individually identifiable health information, the sharing, employment, application, utilization, examination, or analysis of such information within an entity that maintains such information.").
[12] H.R. 847, 117th Cong. (2021); S. 224, 117th Cong. (2021).
[13] S. 3065, 117th Cong. (2021); H.R. 5807, 117th Cong. (2021).
[14] H.R. 6580, 117th Cong. (2022); S. 3572, 117th Cong. (2022).
[15] H.R. 3133, 117th Cong. (2021).
[16] European Data Protection Board, Recommendations 01/2020 on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data, Version 2.0, at 33–34 (June 18, 2021), https://edpb.europa.eu/system/files/2021-06/edpb_recommendations_202001vo.2.0_supplementarymeasurestransferstools_en.pdf; *id.* at 45.

the use of specific PETs—in the form of specific cryptographic protocols or specific algorithmic implementations of particular PET techniques, for example—prematurely. Early adoption of excessively specific standards can stifle innovation in the field of PETs by discouraging openness to new technologies that have developed outside the parameters of such standards; we emphasize again that all aspects of PETs are rapidly evolving, and prematurely articulated standards risk rapidly falling behind the state of the art or, worse still, impeding progress of the state of the art. This situation should be distinguished from that of cryptographic *primitives* underlying many PETs, such as encryption and digital signature schemes, which have been productively standardized; PETs in general are significantly more complex than these basic "building-block" technologies, and as such greater care is required in crafting standards for them.

As such, in the near term, NIST can take a progressively more formal approach to guiding the public regarding the development and use of PETs, beginning with comparatively informal reports and progressing to recommended best practices and guidelines before issuing formal standards. In keeping with our comments elsewhere in this response, we stress that such guidance from NIST would be particularly useful regarding the *optimal selection* of PET technologies for particular types of application, including clear, accessible explanations of the distinct privacy challenges addressed by distinct PETs.

In addition to such guidance and long-term standard-setting efforts from NIST, which can advance development and adoption of PETs at a "macroscopic" level, at a more "microscopic" level short-term prize challenges—spanning a few weeks to a few months—can generate momentum in developing PET-based solutions for specific areas of application and, additionally, raise awareness of PETs in targeted areas of the public or private sectors. Such challenges are particularly useful for bringing together *interdisciplinary* teams of domain experts in the application area and experts in PETs, both from academia and industry. As two examples of productive such prize challenges, we note the iDASH competition addressing privacy-preserving computation and analysis of biomedical data, in particular human genome data,[17] and the United Kingdom Financial Conduct Authority's 2019 Global AML and Financial Crime TechSprint.[18]

Topic 7: *Risks related to PETs adoption*

We have emphasized throughout this response distinct PETs are optimized for distinct use cases, and indeed a given PET may be entirely *inappropriate* for use in a particular use case. Without deliberate efforts to raise public awareness of not only the existence of PETs in general but also the appropriate use of specific PET tools and techniques, unprincipled adoption of PETs poses two significant risks: (1) an inappropriate PET *approach* (*e.g.*, secure multi-party computation, zero-knowledge proof, homomorphic encryption, differential privacy, or federated learning) may

[17] *See* iDASH Privacy & Security Workshop 2022 – Secure Genome Analysis Competition, http://www.humangenomeprivacy.org (last visited July 8, 2022).
[18] *See* 2019 Global AML and Financial Crime TechSprint, https://www.fca.org.uk/events/techsprints/2019-global-aml-and-financial-crime-techsprint (last visited July 8, 2022).

be applied to a particular use case based on a loose understanding that the technology is "privacy preserving"; and (2) even if an application-appropriate PET approach is selected, the user may lack the technical knowledge to understand the privacy guarantees and limitations of a particular *implementation* of that approach, and accordingly may be unable distinguish sound implementations from unsound ones. Our encouragement that NIST play a central role in public guidance regarding PETs is significantly motivated by the need to mitigate these risks.

A related risk is that PETs may be developed or applied without sufficient appreciation or understanding of either the domain of application or the technical aspects of the relevant PETs themselves, particularly as they relate to data privacy and security. Accordingly, it is critical that general policies and specific deployments of PETs be informed by *interdisciplinary* expertise—including, as appropriate, academic experts in the theory underlying PETs, industry practitioners versed in the practical implementation and deployment of PETs, stakeholders with well-informed perspectives on the policy and legal implications of a given use of PETs, and domain experts in the specific area of application. Indeed, given our limited resources as a small technology company focused on the design and development of PETs, it is for this reason that we have declined to opine, for example, on the optimal *substance* of regulations in our comments on Topic 4 above and have simply called for increased regulatory clarity.

Topic 9: *Existing barriers to PETs adoption*

A primary obstacle to the adoption of PETs is inertia due to the overhead—whether economic or logistical—of integrating novel technologies into existing practices and infrastructures. We believe that the benefits of PET adoption at a societal level—and in many cases at the level of individual adopters—outweigh these costs. But the relative technical sophistication required to understand and trust in the benefits of PETs, let alone to adopt them successfully, can skew the cost-benefit analysis away from the adoption of new technological solutions. The Government can help overcome this barrier by demonstrating the feasibility and efficacy of PETs through adoption in its own operations and by encouraging or requiring their use in appropriate situations.

A second barrier to realizing the full potential of PETs is a potentially narrow focus on privacy of *individuals* when considering privacy-preserving technologies, and an attendant focus on informed-consent frameworks and related privacy-assurance mechanisms. To be sure, considerations of individual privacy are vitally important; we simply caution against a narrow conception of privacy that potentially ignores many beneficial applications of PETs. Indeed, PETs can be applied to overcome data-privacy obstacles to collaboration that are based more on competitive or security concerns than on considerations of individual privacy; PETs can enable, for example, joint computation for activities ranging from collaborative risk analysis and mitigation between competitors to industry-wide auctions. Such applications may be less obviously within the ambit of "privacy" than protection of individuals' personal data and accordingly may be neglected in discourse addressing PETs; policies for advancing PETs should be designed to realize the full range of benefits that such technologies can offer.

Thank you for the opportunity to inform the development of a national strategy for privacy-preserving data sharing and analytics and associated initiatives; we hope you find our comments useful.

Sincerely,

Steve Lu
CEO
Stealth Software Technologies, Inc.