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

Confidential Computing Consortium

July 8, 2022

To Whom It May Concern:

Please consider the following submission to the Request for Information on Advancing Privacy-Enhancing Technologies from the Confidential Computing Consortium. The Confidential Computing Consortium (https://confidentialcomputing.io) is a Linux Foundation project "to accelerate the adoption of Trusted Execution Environment (TEE) technologies and standards" and has a diverse membership of hardware and software vendors and cloud service providers (https://confidentialcomputing.io/members/). This response was prepared by the group's Technical Advisory Council with participation from across the membership, and ratified by its Governing Board. The Linux Foundation is a non-profit organization registered in the United States as a 501(c)(6).

The Confidential Computing Consortium has a mandate to engage with governments, standards agencies and regulatory agencies to encourage adoption of Confidential Computing, as well as work with the larger ecosystem and engage with existing and potential end-users of the technologies. It also works with open source projects to further development of implementations. The Confidential Computing Consortium is committed to encouraging open source implementations of Confidential Computing technologies to ensure wide-spread adoption, scalable community involvement, transparency of process, increased security and ease of auditing by relevant interested parties and authorities.

The Confidential Computing Consortium welcomes collaboration with governmental and non-governmental organizations and has mechanisms in place to provide appropriate membership, as well as open technical participation without any membership requirement.

Sincerely,

Stephen R. Walli
Confidential Computing Consortium, Governing Board Chair

Office of Science and Technology Policy
Request for Information on Advancing Privacy-Enhancing Technologies

Response by the Confidential Computing Consortium on Confidential Computing and hardware-based Trusted Execution Environments

Confidential Computing is the protection of data in use by performing computation in a hardware-based Trusted Execution Environment. A technical description of Confidential Computing and the use of Trusted Execution Environments is available in the Consortium's white paper "*A Technical Analysis of Confidential Computing"*: https://confidentialcomputing.io/wp-content/uploads/sites/85/2022/01/CCC-A-Technical-Analysis-of-Confidential-Computing-v1.2.pdf

Information Requested: Respondents may provide information for one or as many topics below as they choose. Through this RFI, OSTP seeks information on potential specific actions that would advance the adoption of PETs in a responsible manner, including on the following topics:

1. *Specific research opportunities to advance PETs:* Information about Federal research opportunities that could be introduced or modified to accelerate the development or adoption of PETs. This includes topics for research, hardware and software development, and educational and training programs. This also includes information about specific techniques and approaches that could be among the most promising technologies in this space.

2. *Specific technical aspects or limitations of PETs:* Information about technical specifics of PETs that have implications for their development or adoption. This includes information about specific PET techniques that are promising, recent or anticipated advances in the theory and practice of PETs, constraints posed by limited data and computational resources, limitations posed by current approaches to de-identification and deanonymization techniques, limitations or tradeoffs posed when considering PETs as well as technical approaches to equity considerations such as fairness-aware machine learning, security considerations based on relevant advances in cryptography or computing architecture, and new or emerging privacy-enhancing techniques. This also includes technical specifications that could improve the benefits or privacy protections, or reduce the risks or costs of adopting PETs.

Confidential Computing is the protection of data in use by performing computation in a hardware-based Trusted Execution Environment (TEE). A comparison with other technologies, including other PETs, is available in Section 4 of "*A Technical Analysis of Confidential Computing*". Key characteristics of TEEs include:

- **Data confidentiality**: Unauthorized entities cannot view data while it is in use within the TEE.
- **Data integrity**: Unauthorized entities cannot add, remove, or alter data while it is in use within the TEE.
- **Code integrity**: Unauthorized entities cannot add, remove, or alter code executing in the TEE.

In addition, implementations may include the following capabilities (also described in "*A Technical Analysis of Confidential Computing")*:

- Code Confidentiality
- Authenticated Launch

- Programmability
- Attestability
- Recoverability

Confidential Computing, by combining these capabilities, can provide two of the three characteristics of the *CIA triad*, Confidentiality and Integrity. It can provide this both for data and code, with the ability to have strong cryptographic assurances derived from attestation measurements and validation. These properties can be used to build applications which are privacy-enhancing, and the deployment of complementary PET software within TEEs may provide extra benefits. There is an existing and growing installed base of computer systems equipped with processors with TEE capabilities, that can support Confidential Computing, including availability in public clouds, both in the United States and globally. The general availability of such systems allows parties with an interest in deploying PETs easy and fast access to their benefits, typically with the ability to support general computation with small performance penalties, rather than having to design and deploy specific algorithms that may operate at significantly slower speeds. The document *A Technical Analysis of Confidential Computing* provides a comparison of some of the characteristics of Confidential Computing and other PETs.

Confidential Computing approaches that support attestation allow cryptographic assurances for every workload or application for each deployment, extending the software supply chain security beyond development and deployment into the runtime environment.

Current research and development in Confidential Computing promises opportunities for TEEs not just in CPUs but in other processors (e.g., GPUs) and the creation of a trust fabric of mutually trusting computational units, with the trust based in cryptographic assurances at runtime.

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. This includes sectors and applications where data are exceptionally decentralized or sensitive, where PETs could unlock insights or services of significant value to the public, where PETs can reduce the risk of unintentional disclosures, where PETs might assist in data portability and interoperability, and sectors and applications where the adoption of PETs might exacerbate risks, including in the areas of privacy, cybersecurity, accuracy of data analysis, equity for underserved communities, and economic competition. This topic covers opportunities to improve the effectiveness of data sharing among specific Federal agencies and between specific Federal agencies and entities outside the Federal Government, including the goals outlined in Section 5 of Executive Order 14058: Transforming Federal Customer Experience and Service Delivery To Rebuild Trust in Government.

Since most TEEs run on general-purpose computing platforms with their associated performance and programming flexibility, the use cases for Confidential Computing span a wide range of sectors. The Confidential Computing Consortium's white paper "*Confidential Computing: Hardware-Based Trusted Execution for Applications and Data"* (https://confidentialcomputing.io/wp-content/uploads/sites/85/2021/03/confidentialcomputing_outreach_whitepaper-8-5x11-1.pdf) contains a section "Use cases for Confidential Computing?" providing some examples including cryptographic key management, Edge and IoT, blockchain, mobile and personal computing devices, public cloud, Point of Sale (PoS) and payment and Multi-Party Computation. These use cases span government, defense and security applications as well as all areas of private enterprise and non-governmental organizations including financial, healthcare, automotive, manufacturing, energy, and beyond. Given the increased global concerns around privacy protection for consumers (and associated regulatory regimes), and the

increasing focus of the cybersecurity industry on protection of data and intellectual property, there are few industries or sectors that would not benefit from Confidential Computing, alongside other PETs.

4. *Specific regulations or authorities that could be used, modified, or introduced to advance PETs:* Information about Federal regulations or authorities that could be used, modified, or introduced to accelerate the development or adoption of PETs. This includes privacy-related rulemaking authorities under the Office of Management and Budget, the Federal Trade Commission, and financial regulatory bodies, as well as acquisition regulations under the Federal Acquisition Regulations. This also includes the Federal authority to set procedures for agencies to ensure the responsible sharing of data. This also covers hiring authorities to recruit Federal employees with expertise to advance PETs, as well as acquisition authorities (e.g., Other Transaction Authority) to procure PETs for development.

There are a number of regulations and industry frameworks that currently lag behind the developments in PETs in general, and Confidential Computing in particular. These include, among others:

- NIST Cybersecurity Framework (https://www.nist.gov/cyberframework), which is widely used within the cybersecurity field. The *Protect Function*, one of the core constituents of the framework, is perfectly suited to benefit from the capabilities provided by Confidential Computing. NIST is currently looking to improve the framework: (https://www.federalregister.gov/documents/2022/02/22/2022-03642/evaluating-and-improving-nist-cybersecurity-resources-the-cybersecurity-framework-and-cybersecurity)
- IEC 62443

We would like to note that threats and regulations evolve at a more rapid pace than rule-making. The most effective laws and regulations are written from a principles- and risk-based perspective, and they should be able to withstand the test of time despite increasingly sophisticated threats and the emergence of new technologies. As an example, the Gramm-Leach-Bliley Act (GLBA) is two decades old, yet the principles of the law remain relevant. However, if rules are written with prescriptive requirements, they are in danger of becoming outdated quickly. This can put security at odds with compliance when security managers need to choose between staying in compliance or using new techniques that are better for the business.

The following are existing regulations and guidance that the financial sector relies on for data handling. Our recommendation is that any recommendations to update guidance principles are objectives-based and avoid prescriptive guidance.

- Cyber Incident Reporting for Critical Infrastructure Act (March 2022)
- Gramm-Leach-Bliley Act (GLBA) Title V, Subtitle A: Disclosure of Nonpublic Personal Information (1999)
- SEC Regulation S-P: Privacy of Consumer Financial Information (November 2000)
- OCC Comptroller's Handbook: Privacy of Consumer Financial Innovation (October 2011)
- FRB Regulation P: Privacy of Consumer Financial Information (December 2001)
- FFIEC Sound Practices to Strengthen Operational Resilience (October 2020)
- FFIEC IT Booklets: Information Security (September 2016)
- FFIEC Guidance: Authentication and Access to Financial Institution Services and Systems (August 2021)
- FFIEC Security in a Cloud Computing Environment (April 2020)
- NIST Special Publication 800-53 (Rev. 5) Security and Privacy Controls for Information Systems and Organizations (September 2020)

5. *Specific laws that could be used, modified, or introduced to advance PETs:* Information about provisions in U.S. Federal law, including implementing regulations, that could be used, modified, or introduced to accelerate the development or adoption of PETs. This includes provisions, safe harbors, and definitions of use, disclosure, safeguards, and breaches. Information may also include comments on how to advance PETs as part of new or proposed legislation, such as that which would create a National Secure Data Service. Information may also include comments on State law or on international law as it applies to data sharing among international entities.

6. *Specific mechanisms, not covered above, that could be used, modified, or introduced to advance PETs:* This includes the development of open-source protocols and technical guidance, the use of public-private partnerships, prize challenges, grants, testbeds, standards, collaborations with foreign countries and nongovernmental entities, the Federal Data Strategy, and data sharing procedures with State, local, tribal, and territorial governments. This also includes interpretations and modifications of standard non-disclosure agreements, confidentiality clauses, data use or sharing agreements, etc.

The Confidential Computing Consortium (https://confidentialcomputing.io) is a Linux Foundation project that "brings together hardware vendors, cloud providers, and software developers to accelerate the adoption of Trusted Execution Environment (TEE) technologies and standards". The effort includes commitments from numerous member organizations and contributions from several open source projects. The Confidential Computing Consortium has a mandate to engage with governments, standards agencies and regulatory agencies to encourage adoption of Confidential Computing, as well as work with the larger ecosystem and engage with existing and potential end-users of the technologies. It also works with open source projects to further development of implementations.

The Confidential Computing Consortium, as part of the Linux Foundation, is committed to encouraging open source implementations of Confidential Computing technologies to ensure wide-spread adoption, scalable community involvement, transparency of process, increased security and ease of auditing by relevant interested parties and authorities. Opportunities for grant-giving to open source projects (either within the Consortium or outside it) should be considered, to encourage commercial implementation and adoption of Confidential Computing.

The Confidential Computing Consortium welcomes collaboration with governmental and non-governmental organizations and has mechanisms in place to provide appropriate membership, as well as open technical participation without any membership requirement.

We believe that appropriate use of Confidential Computing can lower the security risks associated with deployment of smart devices in consumer and industrial scenarios, where insurance may play a role financially. We would thus invite discussion of how insurance regulations might be used to incent use of Confidential Computing technologies in various sectors, including devices ranging from IoT to Edge to public cloud and on-premises computing.

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. This includes risks related to equity for underserved or marginalized groups, the complexity of implementation and resources required for adoption, as well as from conceptual misunderstandings of the technical guarantees provided by PETs. This also includes recommendations on how to measure risk of PETs adoption and conduct risk-benefit analyses of use.

8. *Existing best practices that are helpful for PETs adoption:* Information about U.S. policies that are currently helping facilitate adoption as well as best practices that facilitate responsible adoption. This includes existing policies that support adoption, including in the areas of privacy, cybersecurity, accuracy of data analysis, equity for underserved communities, and economic competition. This also includes information about where and when PETs can be situated within tiered access frameworks for accessing restricted data, ranging from publicly accessible to fully restricted data.

9. *Existing barriers, not covered above, to PETs adoption:* Information about technical, sociotechnical, usability, and socioeconomic barriers that have inhibited wider adoption of PETs, such as a lack of public trust. This includes recommendations on how such barriers could be overcome. Responses that focus on increasing equity for underserved or marginalized groups are especially welcome.

10. *Other information that is relevant to the adoption of PETs:* Information that is relevant to the adoption of PETs that does not fit into any of the topics enumerated above.