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

Deloitte Consulting LLP

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.

Deloitte.

Advancing Privacy-Enhancing Technologies

July 8, 2022

In response to Office of Science and Technology Policy Request for Information

July 8, 2022

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

RE: RFI Response: Privacy Enhancing Technologies

Dear Ms. Hessman:

Deloitte[1] Consulting LLP (Deloitte) is pleased to submit a response to The White House Office of Science and Technology Policy (OSTP) Request for Information (RFI) on Privacy Enhancing Technologies (PETs).

We applaud OSTP's efforts to inform a national strategy on accelerating the development and adoption of PETs. Such technologies have the potential to lead to more trustworthy artificial intelligence systems, while also preserving the privacy of sensitive data.

Despite realizing the importance of privacy protection, policymakers and technical experts alike have struggled to overcome the challenges that current approaches to data privacy have created. Appropriate legislative and regulatory environments may help to reduce some of these frictions, but to be effective they should be complemented by the development and application of new methods for data analysis, such as PETs.

We would be pleased to have an open dialog with OSTP as it continues to explore PETs and reviews responses from this and future RFIs. Please do not hesitate to contact me at (703) 216-4581, should you have any questions.

Sincerely,

Ed Van Buren
Principal

AI in Government Leader
Deloitte Consulting LLP

[1] As used in this document "Deloitte" means Deloitte Consulting LLP. Please see www.deloitte.com/us/about for a detailed description of the legal structure of Deloitte LLP and its subsidiaries. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Table of Contents

Company Profile

Company Name	Deloitte Consulting LLP
Headquarters Location	New York City
Contact Name	Ed Van Buren
Contact Title	Principal, Deloitte Consulting LLP
Contact Email Address	
Contact Phone Number	
Primary Type of Service(s) Provided	Software Development, Professional Services, Management Consulting, Technical Support, Maintenance, and Support Services

Executive Summary

Artificial Intelligence (AI) has been widely recognized as having the potential to transform how services are delivered to citizens and to help government leaders make better informed, data-driven decisions. The National AI Initiative Act represents a significant step in ensuring that the United States continues to lead the world in AI research and development, as well as preparing the nation for the adoption of AI systems across the federal government and, more broadly, society. More recent developments, like the White House Office of Science and Technology Policy (OSTP)'s efforts to engage the public on a "Bill of Rights for an Automated Society," illustrate the importance of aligning data-driven technologies (e.g., AI) with the values that underpin American society.

In preparing a response to this RFI, Deloitte draws upon extensive experience in AI strategy, implementation, and operational projects across the public and private sectors. More than 1,200 Deloitte practitioners are currently engaged with federal, state, and local government clients to harness the power of AI, machine learning (ML), and data analytics. In the commercial sector, Deloitte's comprehensive range of technology solutions and service offerings have helped accelerate AI adoption at the world's leading companies. In addition, Deloitte's AI Institute for Government [1] has served as a prominent hub for AI research, eminence, and applied innovation. Deloitte's Trustworthy AI™ (TAI) framework [2] has also been used in three executive agencies as a blueprint to maximize the benefits of AI, while also effectively managing potential risks and ethical imperatives.

One critical dimension of our TAI framework is privacy protection, a concept that has also been identified by the federal government as a strategic pillar of trustworthy AI systems. Despite realizing the importance of privacy protection, policymakers and technical experts alike have struggled to overcome the challenges that current approaches to data privacy have created. For example, parties that want to collaborate on building AI systems using shared data must first negotiate complex data use agreements. These data use agreements often take many weeks or months to finalize, creating delays and limiting the ability of organizations to fully harness the power of AI.

Appropriate legislative and regulatory environments may help to reduce some of these frictions, but these efforts should be complemented by the development and application of new methods for data analysis, such as privacy enhancing technologies (PETs). PETs refer to a broad set of algorithms, software, and hardware tools that preserve the privacy of sensitive data during analysis. Many of these technologies, like homomorphic encryption and zero knowledge proofs, have been researched and applied extensively in the fields of mathematical cryptography and computer science since the early 1980s. Other PETs, such as federated learning and differential privacy, were formalized more recently to help securely analyze the massive amounts of sensitive digital data generated from smartphones and wearable electronics.

We believe that successful use of PETs for trustworthy AI will require the federal government to take a holistic, "systems-level" perspective. Proper application of PETs will involve an understanding of the interconnections between AI and a host of related topics, including digital identity, the Internet of Things (IoT), and quantum computing. Moreover, by treating PETs as modular components that can be integrated into existing information systems, the government can lower some of the technical barriers to adoption. Identifying opportunities to combine different PETs into the same application or technology stack could also help to increase data privacy protections.

One notable PET that can serve as a useful template for integrated, privacy preserving data analytics systems is federated learning (FL). FL refers to a setting in which machine learning models are trained collaboratively by a network of participants using a distributed network of datasets that are independently owned and maintained. In a FL system, no raw source data is explicitly shared between the collaborators.

FL can be applied to address important challenges in key areas, such as public health, national security, and civil infrastructure. Using data analytics and AI to improve disease surveillance, combat financial crimes, enhance military interoperability at the tactical edge, and build the next generation of "smart" cities will require a whole-of-society response. Different government agencies, departments, and even private sector parties can use FL to collaborate on critical data analytics efforts, while also mitigating risks to privacy.

To realize the full potential of PETs in creating trustworthy AI systems, the federal government can focus on the mechanisms that are proven to support scientific and technological innovation. Among these are workforce training and educational programs, incubation of public-private partnerships, and the development of common technical standards. By encouraging public-sector workforce training programs to integrate PETs into their curriculums, OSTP can help lead with fellow federal partners to ensure the next generation of technical practitioners and administrators are well equipped to apply PETs effectively. Similarly, by expanding the National Science Foundation (NSF)-led National AI Research Institutes to include PETs as a research area, the federal government can establish adequate funding for these important technologies.

An effective national strategy to accelerate the adoption of PETs should also address the key roles that external non-government partners can help play in implementing these technologies effectively. By creating an ecosystem that incubates relationships between academia, early-stage start-ups, and leading technology companies, the federal government can build on the collective strengths of a robust group of stakeholders. Low-risk collaboration opportunities, like agency-specific technology sprints, public competitions, and state-wide challenges can help to create a robust stakeholder ecosystem. In turn, this ecosystem can help shape common standards for how PETs are designed, developed, and implemented as part of more complex data analytics systems. As a leader in both AI and private-sector consulting, Deloitte has deep experience in developing commercial technology ecosystems, managing alliance partnerships, and fostering productive cross-sector collaboration on AI and data analytics.

Deloitte stands ready to support OSTP in creating a national strategy to advance PETs. We are confident in the vision put forth by Dr. Alondra Nelson for "a world where our technologies reflect our values and innovation opens the door to solutions that make us more secure." [3] We believe that such a future, designed with our society's core values in mind, is well within reach.

Introduction

OSTP's inquiry into PETs marks a significant milestone for the government in its efforts to advance technological innovation that adheres to our society's core values. In support of OSTP's efforts, we are pleased to share our knowledge from our historical and ongoing experience advising public sector and commercial organizations on the implementation of data analytics, AI, and privacy solutions.

In this document, we focus on three topics: specific technical aspects and limitations of PETs (RFI Topic 2), sectors that would benefit from the adoption of PETs (RFI Topic 3), and specific mechanisms that could be used, modified, or introduced to advance PETs (RFI Topic 6). The scope of our response in the first two sections is restricted to a detailed discussion of federated learning (FL), a PET with great potential for rapid adoption across the public and private sectors. While these three topics are the focus of this response, Deloitte's expertise spans across a range of related technical and policy areas. We would be happy to offer our perspective on additional topics in another forum.

Response to Topic 2

Specific Technical Aspects and Limitations of Federated Learning

In this section, we begin with an overview of the technical aspects of federated learning. Next, we illustrate the aspects of FL *systems*, as they pertain to driving adoption of multiple PETs. We then outline the technical limitations and operational risks of FL systems and conclude with a discussion of recent advancements which help to mitigate some of the risks of adopting FL systems.

Federated Learning enables privacy-preserving collaborative analytics

FL is a form of multi-party machine learning that has been used in the consumer technology sector since 2017. FL techniques have been applied to train text-prediction models, improve voice-recognition models, and build other machine learning models in settings where training data could not be centralized due to privacy concerns and bandwidth limitations [4] [5]. FL allows for a distributed network of data providers to collaboratively train a global machine learning model without revealing any of the underlying training data. Each participant performs a learning or analytics task on its own local (private) dataset, and only shares information about the result of that task, rather than the underlying dataset itself.

This approach to machine learning presents clear opportunities to preserve data privacy by bypassing the direct access or transfer of raw data. As a result, organizations can choose FL as a strategy to collaborate on machine learning tasks even when privacy, security, or trust concerns prevent data centralization. The ability of FL to effectively expand the scope of data available for mission-critical AI systems makes it especially relevant in the public sector. Furthermore, when organizations find the same type of data spread across many locations or providers, FL can be used to build models which are more robust, generalizable, and representative. By allowing data providers from different locations and organizations to collaborate in a privacy-preserving way, FL can serve a key role in promoting equitable benefits of AI for participants that may not have access to large volumes of high-quality data.

It is important to note that the key principles that underpin FL can be extended to create more general purpose "federated analytics" systems. Such systems can be used for arbitrary data analysis tasks, like calculating metrics across a distributed network of datasets.

Federated Learning can serve as a vehicle for broader PET adoption

Like all forms of machine learning and data analysis, FL is not applied in a vacuum. FL models are trained, tested, and deployed as part of a complex *system*. Organizations that wish to use FL for collaborative analytics must design and develop an integrated system that involves data storage mechanisms, secure messaging and communication protocols, and distributed compute capabilities. Each of these system functions will involve processing sensitive data, creating a need for security and privacy protections. These privacy and security needs, in turn, can be addressed with other PETs. Hence, an FL system designed for the purpose of high-security, privacy-preserving collaborative data analytics can serve as a platform to integrate a broader set of PETs [6].

Limitations of Federated Learning

- **FL model development can involve significant technical "overhead":** Performing a ML or analytics workflow in an FL setting requires FL-specific aggregation and optimization algorithms suitable for the desired task [7, p. §3]. In many cases, FL involves system-level parameter tuning, code-refactoring, and new algorithm development, leading to longer development cycles and a need for specialized technical expertise.
    - *Advancements:* Technical overhead can be reduced by relying on open research and open-source tools. Researchers are developing and providing code for algorithms which are tailored to specific settings or needs [8] [9], or to support new families of models [10]. At the same time, popular toolkits for FL systems have common algorithms supported out-of-the-box [11] and are increasingly focused on streamlining the developer experience for implementing arbitrary workflows [12].
- **FL model performance can be degraded by highly variable data:** The accuracy, generalizability, and overall performance of FL models are heavily influenced by the quality, heterogeneity, and structure of training data. Different organizations may sample different populations, employ different data collection methods, and have different quality control resources available. Ensuring that an aggregate model will converge with high performance on data distributions which are not necessarily independent and identically distributed (non-IID) is a major challenge in FL.
    - *Advancements:* Strategies to improve FL performance on non-IID data have emerged in various settings including implementing new FL optimization algorithms [13], relying on synthetic data generation [14], augmenting FL workflows [15], or even redefining the target FL workflow to prioritize personalized models over a global model [16].
- **Heterogeneous computing environments can make it difficult for participants to coordinate:** Since FL models are trained collaboratively by a distributed network of participants, it is essential that there are common operating standards across the network. Each participant must have machines with adequate compute power, sufficient data storage capacity, and reliable communication channels. Depending on the setting, there may be significant effort needed to configure computer systems consistently across different participants.
    - *Advancements:* The rapid growth in adoption of containerization tools like Docker is continually making it easier to standardize software environments on different machines [15]. In addition, newly released FL software toolkits are making use of well-established protocols such as gRPC to simplify communication between participating machines. Discrepancies in compute power and availability of participants is also being addressed by researchers exploring techniques of asynchronous FL, active sampling, and fault tolerance [18].

Privacy and Security Risks associated with Federated Learning

- **Model training can be undermined by malicious participants:** The inability to directly inspect the training data from each participant leaves FL models susceptible to data poisoning attacks. These attacks occur when the private datasets or updates returned during federated model training are crafted by a malicious participant to negatively impact the model's quality. Such attacks may come from a single malicious participant or a coordinated group of participants forming a Sybil Attack [19]. If not properly addressed, these attacks can be used to degrade model performance or embed backdoors to make future adversarial attacks more effective.
- **Sensitive information can be reverse engineered:** Models created during federated learning are usually shared between participants during a learning task. In the academic literature, it has been shown to be possible for

participants to perform data reconstruction or membership inference attacks on these models [20]. Without sufficiently strong protections in place, such attacks make FL systems vulnerable to leaking sensitive data.

Risk Mitigation Strategies

Well-designed FL systems typically have strong data privacy requirements, providing an opportunity for the integration of complementary PETs. This integration not only provides the privacy guarantees of each individual PET but also serves to mitigate some of the risks of FL outlined above. Below, we present examples of PETs that have been used to further enhance the privacy protections of a FL system:

- ❑ **Differential Privacy** can be applied to participant information before it is shared so that the contributions have additional privacy guarantees, mitigating the effectiveness of data reconstruction or membership inference attacks [5].
- ❑ **Homomorphic Encryption** can be applied to participant information before it is shared so that the contributions have additional security guarantees, allowing for all participant contributions to remain encrypted throughout the workflows [4].
- ❑ **Secure Aggregation** can be used to enforce cryptographically secure aggregation of participant information so that no contribution can be inspected individually, preventing attacks from targeting specific participants [2].
- ❑ **Zero Knowledge Proofs** can be used to defend against poisoning attacks hidden in participant contributions without violating secure aggregation or revealing participant information [17].
- ❑ **Private Set Intersection** can be used in FL settings to perform entity resolution between participant datasets in a privacy-preserving manner [18].
- ❑ **Trusted Execution Environments** can provide confidentiality and integrity guarantees on local data processing for participants in a FL system [20].

Since FL system architects often turn to other PETs to address a spectrum of privacy requirements and security considerations, it is likely that encouraging greater adoption of FL systems will in turn increase understanding, adoption, and advancement of PETs more generally.

Response to Topic 3

Sectors that would benefit from adoption of Federated Learning

In this section, we seek to direct OSTP’s attention to the key areas in which FL systems are especially useful. Data analytics technologies have long been utilized by public and private sector organizations, but increasing desires for consumer privacy and control, coupled with a growing body of privacy regulation, means that the need for managing risk is growing more complex every day. We believe that adoption of FL systems has the potential to drive transformative change across many sectors, including healthcare, national security, finance, and civil infrastructure.

Healthcare

Policymakers and medical practitioners alike have long envisioned a public health system focused on predicting and proactively preventing illnesses, rather than providing services retroactively. Realizing this vision will require a diverse group of public and private sector partners working towards a unified goal of sharing data-driven insights with one another in real time. The future of public health will depend on an ability to fuse real-time data from multiple sources to empower health leaders to identify health trends and respond to potential threats. In such a future, partners can both contribute to and benefit from predictive data analytics.

FL systems and other PETs would allow for patient information to stay protected on the de-centralized data silos in which they are usually stored, while also allowing for collaboration between hospitals, insurers, researchers, and health departments on critical public health issues.
In 2021, an applied research test used data from a global network of institutes to train a FL model that predicted future

oxygen requirements of symptomatic patients with COVID-19 using inputs of vital signs, laboratory data, and chest X-rays [24]. The model achieved an average area under the curve (AUC) > 0.92 for predicting outcomes at 24 and 72 h from the time of initial presentation to the emergency room. Additionally, it provided 16% improvement in average AUC measured across all participating sites and an average increase in generalizability of 38% when compared with models trained at a single site using that site's data. This study showed that FL resulted in collaboration without loss of identifiable data and produced a result that generalized across heterogeneous, decentralized datasets for prediction in patients with COVID-19.

Project MELLODDY, which concluded in May 2022 after three years of study, is another example of how FL has been successfully applied in the healthcare sector [25]. In this project, pharmaceutical companies collaborated to find a way to combine resources and research – without the loss of intellectual property (e.g. data) – to facilitate the development of new biomedical drugs. Using a FL system, these companies collaborated to improve performance of predictive models and reduce decision time needed to research next steps in drug discovery.

National Security

The benefits of a fully integrated FL system are already known to the Department of Defense (DoD). A 2021 call for industry response (Opportunity ID: TW-22-0003) was issued to address complications associated with integrating ML and AI in assets with Joint All-Domain Command and Control (JADCC). US Naval fleets are frequently inconsistent in both hardware and software, due to the changing landscape of asset manufacturers, operators, contracts, and department policy. Consequently, assets may have different network bandwidths, memory limits, and local data distributions. These are problems addressed by a FL system, as learning can occur asynchronously and independently between node machines. Base requirements for learning are configurable between Naval assets and allow for heterogeneous data sources, heterogeneous hardware bandwidth limitations, and secured communication between assets of JADCC. The result of building an integrated FL system could enable a greater degree of interoperability to increase mission capabilities.

FL systems would benefit national security beyond the tactical edge, as well. In March 2021, the Biden administration issued an executive order identifying supply chain reforms as critical to protecting US economic and national security interests [26]. Current supply chain management is set up on point transactions between top-level suppliers and buyers. The vulnerable upstream supply chain network is more opaque, largely due to hesitation by organizations to share information that could compromise competitive position, reveal compliance posture, or highlight security concerns. FL systems could allow for analysis directly on siloed data that can never be pooled directly due to concerns over privacy, intellectual property, and sovereignty. The result could be a federated global supply chain model built on top of a network of protected data, giving the federal government a more accurate and real-time assessment of supply chain strength and resilience.

Financial Crime Compliance

Private financial institutions sit at the front line of defense against financial crimes by detecting and flagging transactions based on suspicious activity. However, this process is costly and time-intensive due to the need for manual validation. A recent study sought to provide support by developing a FL system with additional PETs (i.e., homomorphic encryption and differential privacy) to increase model accuracy of detecting suspicious activity [27]. Leveraging this FL system, researchers showed accuracy improvements by as much as 20% in comparison to individual financial institutions. These results suggest that private financial entities can use FL systems to collaboratively investigate financial crimes without exposing private data.

FL systems can also be applied in similar settings to detect sanctions evasion, tax fraud, and other financial crimes without compromising the privacy of individual citizens or companies. The potential for PETs like FL to help investigators detect these types of financial crimes has been clearly recognized by the White House, which recently announced PETs-focused prize challenges involving the U.S. Financial Crimes Enforcement Network [3].

Civil Infrastructure

The bipartisan $1.2 trillion Infrastructure Investment and Jobs Act (IIJA), signed by President Biden in November 2021, indicates a clear national priority that both Congress and the President support [28]. With this infusion of support and funding to update and modernize civil infrastructure, the need for PETs to protect citizens' privacy and ensure efficient access to national resources (i.e., power, water, transportation, internet) is paramount.

Over 60% of infrastructure officials – local and federal – who responded to Deloitte's Future of Infrastructure survey identified AI as having the largest impact on infrastructure plans over the next three years [29]. Today, even physical infrastructure increasingly entails a digital component, evidenced by recent cyberattacks against critical infrastructure operators, schools, hospitals, manufacturers, banks, and others [30]. Additionally, as autonomous vehicles become more common, they will require compatibility with vehicle-to-vehicle, vehicle-to-grid, and Internet of Things (IoT) requirements, all of which come with an increased need for controlled learning and communication so that privacy considerations can be identified and protected. PETs, like FL, can play an important role in the development of new "smart" infrastructure, as communication and learning will be vital as new technologies are adopted with privacy in mind.

Among other provisions, the IIJA directs the Environmental Protection Agency to establish grant funding programs for the purpose of reducing cybersecurity vulnerabilities. In addition, it contains provisions related to high-performance computing and smart manufacturing [28]. This funding represents a potential opportunity to pursue the advancement of FL in identifying cybersecurity threats and facilitating smart manufacturing capabilities. In alignment with White House goals for transparency and improving citizens' access to services and service delivery experience [31], FL systems could address rising concerns surrounding security, access, and efficiency of critical infrastructure.

Response to Topic 6

Specific mechanisms that could be used, modified, or introduced to advance PETs

FL systems and other PETs can only succeed in a well-designed and efficiently functioning ecosystem that serves to improve the lives of US citizens. Such an ecosystem should focus on stakeholder organization and communication to prioritize transparency, privacy, and security. In this section, we describe some core mechanisms that could be introduced, modified, or utilized to achieve the privacy preserving and collaborative analytical benefits of FL systems and other PETs.

Integrate PETs into education and workforce training programs

Public sector practitioners should have extensive technical training on PETs to ensure they can successfully apply them at work. Agency leaders also need to understand PETs and their strategic applications for cross collaboration that range in requirements and feasibility. This can be achieved through integrating PET curriculums into existing educational and training programs.

For example, the Public Health Informatics and Technology (PHIT) Workforce Development Program could include additional technical training on PETs for future system engineers. Existing programs, such as National Institute of Standards and Technology (NIST)'s Differential Privacy Program, could be modified to increase PET education. At the academic level, scholarship programs for graduate studies could be tailored to PET research and include employment opportunities. For example, the National Science Foundation (NSF) AI Institutes could also be expanded to support PET researchers. Academic-driven initiatives would also enable students – America's future workforce – to experience the possibility of careers in the science, technology, engineering, and mathematics (STEM) fields. This could ensure a future pipeline of workers from diverse backgrounds who are both technically prepared and privacy-aware.

Encourage innovation and organic partnerships through a stakeholder ecosystem

PETs are an expanding field of privacy and security methods that benefit from the latest techniques to remain up to date and effective. Decision makers in the public sectors should recognize that collaborating with the private sector on PET development could help avoid strategic gaps in knowledge or capabilities. Encouraging innovation, maintaining strong partnerships, and referencing open-source tools are critical for this purpose. By leveraging agency-specific "tech sprints," public competitions, and managing procurements through dedicated Other Transaction Authorities (OTAs), the federal government can help foster innovation.

By replicating the Department of Veterans Affairs' (VA) National AI Institute's AI Tech Sprints across other agencies, the federal government could provide low-risk testbeds for PETs in the public sector. Successful outcomes could then be followed by rigorous testing and evaluation for production use. Similarly, current programs such as SBIR Catalyst and the Growth Accelerator Fund could provide a pipeline for constant advancement in implementation and security, and government focused venture firms such as In-Q-Tel or NSF's America's Seed Fund could further ideation development. Internally, establishing research partnerships amongst academic and research institutions provides necessary advisory reference and quality assurance support.

To ensure strong partnerships, we encourage the government to provide select smaller private entities access to its own large data resources, when appropriate. When private institutions access large pools of public data to train their own models, they can expedite results and make potentially large impacts. For example, many private health institutions might have the expertise to develop advanced models but lack robust data to train their models. Access to robust and plentiful healthcare data collected by the government could expedite healthcare research and discovery.

Develop a standardized set of best practices for collaboration

We see the value in having a foundational set of standardized best practices to foster innovative ideas, talent, and resources. For instance, the NIST developed its Cyber Security Framework (CSF), consisting of standards, guidelines, and best practices to manage cybersecurity risk. The CSF has proved effective at managing a cohesive set of standards when implementing cyber security programs. Establishing similar guidelines to the CSF, among others [32], could provide the framework and foundation necessary to develop, use, and evaluate FL and PET products, systems, and services. When Deloitte supported NIST to develop the CSF, we conducted interviews with the private sector on NIST's behalf and one of the most important aspects of the CSF that was highlighted was the flexibility to achieve the desired outcomes in a number of pre-existing ways.

We also helped to establish NIST's Privacy Engineering Program (PEP), which includes an operating framework for contributing relevant tools and use cases. Such a framework regarding privacy techniques and standard practices could provide a template for PET pilot programs. Additional considerations, such as developing a common set of approved software tools, could potentially mitigate privacy and security risks—as well as help to uncover any unanticipated consequences of normally-operating PETs.

Focus on transparent and clear communication about PETs

Much of the data used to improve the daily lives of US citizens contains sensitive and private individual information. As such, it is vital to provide assurances and instill accountability when utilizing individual-level data. We recommend internal awareness communiques regarding PET impact to be circulated among government practitioners and leadership. These communications could take the form of newsletter and blogs, like the NIST's PEP Program. Externally, easily explainable public awareness campaigns can help highlight the use cases and benefits of PETs. Consistent communication can help reinforce transparency and ensure that safeguards remain vigorous and effective.

Conclusion

Deloitte is excited about a future that is enhanced through data analytics and AI, especially when the appropriate privacy protections are identified and integrated. PETs have the potential to enable our society to operate more safely and efficiently, while also protecting individual liberties and privacies. Across a broad swathe of sectors, PETs can be used to build more trustworthy and secure data analytics systems. It is with this promise in mind that we look forward to supporting OSTP and other federal government partners in establishing a national strategy to accelerate the development and adoption of PETs.

References

[1] "Deloitte AI Institute for Government," Deloitte US, [Online]. Available: https://www2.deloitte.com/us/en/pages/public-sector/articles/artificial-intelligence-government-sector.html.

[2] "Trustworthy Artificial Intelligence (AI)™," Deloitte US, [Online]. Available: https://www2.deloitte.com/us/en/pages/deloitte-analytics/solutions/ethics-of-ai-framework.html.

[3] U.S. White House Briefing Room, "U.S. and U.K. Governments Collaborate on Prize Challenges to Accelerate Development and Adoption of Privacy-Enhancing Technologies," Press Release, 13 June 2022. [Online]. Available: https://www.whitehouse.gov/ostp/news-updates/2022/06/13/u-s-and-uk-governments-collaborate-on-prize-challenges-to-accelerate-development-and-adoption-of-privacy-enhancing-technologies/.

[4] B. McMahan and D. Ramage, "Federated Learning: Collaborative Machine Learning without Centralized Training Data," Google AI Blog, April 2017. [Online]. Available: https://ai.googleblog.com/2017/04/federated-learning-collaborative.html.

[5] K. Hao, "How Apple personalizes Siri without hoovering up your data," MIT Technology Review, December 2019. [Online]. Available: https://www.technologyreview.com/2019/12/11/131629/apple-ai-personalizes-siri-federated-learning/.

[6] D. Stewart, A. Bucaille and G. Crossan, "Homomorphic encryption and federated learning can underpin more private, secure AI," Deloitte Insights, December 2021. [Online]. Available: https://www2.deloitte.com/xe/en/insights/industry/technology/technology-media-and-telecom-predictions/2022/homomorphic-encryption-federated-learning.html.

[7] P. Kairouz, B. McMahan and et al., "Advances and Open Problems in Federated Learning," 2019. [Online]. Available: https://arxiv.org/abs/1912.04977.

[8] S. Reddi, Z. Charles, M. Zaheer and et al., "Adaptive Federated Optimization," 2020. [Online]. Available: https://arxiv.org/abs/2003.00295.

[9] T. Li, M. Sanjabi, A. Beirami and V. Smith, "Fair Resource Allocation in Federated Learning," 2019. [Online]. Available: https://arxiv.org/abs/1905.10497.

[10] Y. Wu, S. Cai, X. Xiao and et al., "Privacy Preserving Vertical Federated Learning for Tree-based Models," 2020. [Online]. Available: https://arxiv.org/abs/2008.06170.

[11] "NVIDIA FLARE," [Online]. Available: https://developer.nvidia.com/flare.

[12] "Flower: A Friendly Federated Learning Framework," [Online]. Available: https://flower.dev/.

[13] L. Zhang, Y. Luo, Y. Bai and L.-Y. Duan, "Federated Learning for Non-IID Data via Unified Feature Learning and Optimization Objective Alignment," 2021. [Online]. Available: https://openaccess.thecvf.com/content/ICCV2021/papers/Zhang_Federated_Learning_for_Non-IID_Data_via_Unified_Feature_Learning_and_ICCV_2021_paper.pdf.

[14] H. Chen and H. Vikalo, "Federated Learning in Non-IID Settings Aided by Differentially Private Synthetic Data," 2022. [Online]. Available: https://arxiv.org/abs/2206.00686.

[15] Y. Zhao, M. Li, L. Lai and et al., "Federated Learning with Non-IID Data," 2018. [Online]. Available: https://arxiv.org/abs/1806.00582.

[16] Y. Huang, L. Chu, Z. Zhou and et al., "Personalized Cross-Silo Federated Learning on Non-IID Data," 2021. [Online]. Available: https://www.aaai.org/AAAI21Papers/AAAI-5802.HuangY.pdf.

[17] "Docker," [Online]. Available: https://www.docker.com/.

[18] T. Li, A. Kumar Sahu, A. Talwalkar and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," 2019. [Online]. Available: https://arxiv.org/abs/1908.07873.

[19] C. Fung, C. J. Yoon and I. Beschastnikh, "The Limitations of Federated Learning in Sybil Settings," 2020. [Online]. Available: https://www.usenix.org/system/files/raid20-fung.pdf.

[20] V. Mothukuri, R. M. Parizi, S. Pouriyeh and et al., "A survey on security and privacy of federated learning," 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167739X20329848.

[21] T. Nguyen and M. T. Thai, "Preserving Privacy and Security in Federated Learning," 2022. [Online]. Available: https://arxiv.org/abs/2202.03402.

[22] L. Lu and N. Ding, "Multi-party Private Set Intersection in Vertical Federated Learning," 2020. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9343209.

[23] F. Mo, H. Haddadi, K. Katevas, E. Marin, D. Perino and N. Kourtellis, "PPFL: Privacy-preserving Federated Learning with Trusted Execution Environments," 2021. [Online]. Available: https://arxiv.org/abs/2104.14380.

[24] I. Dayan, H. R. Roth, A. Zhong and et al., "Federated learning for predicting clinical outcomes in patients with COVID-19," 2021. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/34526699/.

[25] Project MELLODDY, [Online]. Available: https://www.melloddy.eu/.

[26] Executive Office of the President, "Executive Order 14017: America's Supply Chains," 24 February 2021. [Online]. Available: https://www.federalregister.gov/documents/2021/03/01/2021-04280/americas-supply-chains.

[27] T. Suzumura, Y. Zhou, N. Baracaldo, G. Ye and et al., "Towards Federated Graph Learning for Collaborative," IBM Research, 2019. [Online]. Available: https://arxiv.org/pdf/1909.12946.pdf.

[28] 117th Congress, "H.R. 3684 - Infrastructure Investment and Jobs Act," 2021-2022. [Online]. Available: https://www.congress.gov/bill/117th-congress/house-bill/3684/text.

[29] Deloitte Center for Government Insights, "The future of United States infrastructure: A survey of infrastructure trends," Deloitte Insights, 2022. [Online]. Available: https://www2.deloitte.com/us/en/insights/industry/public-sector/future-of-infrastructure.html.

[30] U.S. White House Briefing Room, "Press Briefing by Press Secretary Jen Psaki and Deputy NSA for Cyber and Emerging Technologies Anne Neuberger, March 21, 2022," Press Briefings, 21 March 2022. [Online]. Available: https://www.whitehouse.gov/briefing-room/press-briefings/2022/03/21/press-briefing-by-press-secretary-jen-psaki-and-deputy-nsa-for-cyber-and-emerging-technologies-anne-neuberger-march-21-2022/.

[31] Executive Office of the President, "Executive Order 14058: Transforming Federal Customer Experience and Service Delivery To Rebuild Trust in Government," 13 December 2021. [Online]. Available: https://www.federalregister.gov/documents/2021/12/16/2021-27380/transforming-federal-customer-experience-and-service-delivery-to-rebuild-trust-in-government.

[32] National Institute of Standards and Technology, "AI Risk Management Framework: Initial Draft," 17 March 2022. [Online]. Available: https://www.nist.gov/system/files/documents/2022/03/17/AI-RMF-1stdraft.pdf.