NSF Convergence Accelerator: The Future of Privacy Technology (C-Accel 1939288)

Jules Polonetsky, CEO, Future of Privacy Forum
Jeremy Greenberg, Policy Fellow, Future of Privacy Forum

Background:

Data is the lifeblood of modern organizations, and data governance is in many ways *the* modern business and societal problem. The impact of data-intensive technologies is increasing as the pace of innovation accelerates with resultant privacy risks from more extensive collection and processing of personal data. The impact (and need) reaches across every domain and every sector of the economy: business and industry, government and higher education, foundations, and civil society organizations are all looking for ways to embed privacy and data protection in their operations and get access to the tools they need to safeguard personal data.

The technical, organizational, and logistical complexity of modern-day data governance and the need to protect and respect individual privacy rights constitute a "grand challenge problem" which will require a grand challenge-level focus and investment to solve. The Future of Privacy Forum (FPF) has proposed a privacy technology track for NSF's Convergence Accelerator to promote industry and academic collaboration on technical solutions to protect privacy.

Initial Project and Visioning Outline:

Interaction between industry and academic researchers is a critical element to what will be achieved in advancing privacy technology, helping privacy researchers understand technological and business needs and developments on the ground.

With NSF's support[1], FPF undertook an initial project to identify future directions and requirements of privacy technology from the industry perspective. A survey was designed and administered to industry privacy leaders to provide input on the future of privacy technology as relates to their business and policy needs and objectives.[2]

In our summary and analysis as follows, we have identified three major areas that are especially ripe for investment and development: 1. Privacy Enhancing Technical Tools; 2. Administrative and Compliance Tools; and 3. Self-Regulatory and Policy Tools. These are discussed in detail below

---

[1] This material is based upon work supported by the National Science Foundation under Grant No. 1939288.

[2] The survey was presented and discussed with industry privacy leaders who are represented on FPF's Advisory Board at three meetings in Washington, DC; New York City; Los Angeles in October and November 2019. The survey instrument is provided in the appendix.

to provide guidance to the market, to the academic community, and particularly to NSF in selection and advancement of the 2020 Convergence Accelerator.

1. Privacy Enhancing Technical Tools

a. **De-identification tools**: De-identification is a process or technique for removing direct, and known indirect, identifiers from data sets. De-identification balances the need for strong individual privacy protection with beneficial uses of data, such as medical research and emergency response. However, as long as data can be utilized, there is always a risk that data can be re-identified and linked back to the individuals from which the data was derived.[3] De-identification tools can help preserve both privacy and utility by providing methods for de-identifying data, such as injecting noise into a data set, sharing de-identified data through secure methods like differential privacy, and measuring key statistics related to the utility, privacy, and provability of the data. Although scientists, lawyers, and regulators often refer to de-identification as a clearly defined standard for safeguarding privacy, de-identification operates across a spectrum of identifiability balanced with the utility of the data.[4] Moreover, de-identification is required for compliance with a number of federal laws in areas such as: education (FERPA), health (HIPAA), and aviation (CFR Article 49 - Transportation), among others.

Two of the most prominent models of de-identification are: 1. Privacy Preserving Data Mining (PPDM); and 2. Privacy Preserving Data Publishing (PPDP). PPDM refers to techniques for maintaining individual privacy regarding sensitive personal information used to populate public statistics, such as the results of a survey involving confidential information. One of the most critical techniques of PPDM is differential privacy. Differential privacy preserves confidentiality through the injection of "statistical noise" into the values of a data set prior to public release. The injection of noise works to prevent the identification of any specific individual, while still using data collected from individuals. Differential privacy is an ongoing area of research and is applied to an increasing number of use cases, such as the 2020 Census, web browsers gathering statistics regarding user processes, and creation of synthetic data for machine learning.

PPDP permits the use of private data by outside researchers to conduct critical studies, while maintaining strong privacy for individuals from whom the data was derived. One example of a PPDP technique is synthetic data generation, which either creates a data set consisting of some or all "synthetic" data similar to actual individual data that cannot easily be matched to the original data subject.

In addition to de-identification techniques, companies looking to utilize data while maintaining individual privacy often use re-identification tools and other techniques to test the level of risk of re-identification of de-identified data sets.

---

[3] Simson L. Garfinkel, *NISTIR 8053: De-Identification of Personal Information* (Oct. 2015) at 1, *available at*: https://nvlpubs.nist.gov/nistpubs/ir/2015/NIST.IR.8053.pdf.
[4] *See* FPF, *A Visual Guide to Practical Data De-Identification* (Apr. 25, 2016), *available at*: https://fpf.org/2016/04/25/a-visual-guide-to-practical-data-de-identification/.

Our initial survey included conversations at three meetings with privacy leads from a total of 40 companies. In each meeting, companies indicated significant interest in de-identification, but in most cases had limited awareness of the academic state-of-the-art. To the extent companies were versed in the basics of differential privacy or homomorphic encryption, they were unclear which use cases would benefit from these technologies or had the impression that easy to use implementations were not available.

b. **Harm mitigation tools:** Organizations are increasingly pursuing strategies to mitigate potential harms arising from use of data. These technical, legal, and policy tools are particularly useful when data minimization or de-identification techniques are not practical. Harm mitigation technologies have reduced spam email, combated phishing attacks, and mitigated the risks of financial fraud. Emerging tools promise to decrease unwanted robocalls, promote account security, combat the spread of non-consensual intimate images, and reduce discriminatory or unfair algorithmic decisions.

According to the Federal Communications Commission (FCC), robocalls are the most frequent consumer complaint to the agency by a wide-margin. Some studies suggest that U.S. consumers collectively received a total of 4 billion robocalls per month in 2018.[5] Robocalls are phone calls that use an autodialer to deliver a recorded message to a consumer's telephone number. Relatedly, caller ID spoofing occurs when a caller "spoofs" another number, often making it appear to be an incoming call from a local number or a number normally trusted by a consumer. Robocalls and spoofs result in harms ranging from user inconvenience to malicious scams defrauding consumers out of their financial assets, while collecting sensitive personal information and harming user privacy. Mitigation tools in development by the major mobile carriers and other third party security firms, including authentication tools, behavioral analysis and tracking of unusual calling, and call blacklisting are necessary to mitigate the increasingly significant technology used by bad actors to scam consumers.

Promoting user account security is increasingly top-of-mind throughout the technology landscape in the wake of several notable data breaches involving consumers' financial, location, health, and other sensitive data. In addition to consumer privacy harms, poor account security can pose health, public safety, and economic risks. Because it is unrealistic to rely on consumers to possess the technical knowhow and compunction to maintain strong account security (and perhaps almost as unlikely for SME's to have the internal resources to promote strong account security), industry is increasingly dependent on up-to-date, and affordable, account security tools needed to mitigate both rudimentary and sophisticated security threats.

Along with robocalls and account cyberthreats, the creation and sharing of non-consensual images online has reached epidemic proportions. 49 states along with Washington, DC have responded with "revenge porn" laws, while Congress considers solutions at the federal level, including revisions to § 230 of the Communications Act to mitigate associated harms. While many

[5] FCC, *The FCC's Push to Combat Robocalls and Spoofing* (last visited Dec. 27th, 2019, 8:43PM), *available at*: https://www.fcc.gov/about-fcc/fcc-initiatives/fccs-push-combat-robocalls-spoofing.

state laws provide strong remedies for the victims of these unique privacy and security harms resulting from the non-consensual image creation and sharing, the tools used to generate these images are often open source and accessible to amateurs, while detection of non-consensual and fake images proves sometimes difficult for experts. Tools for the detection, removal, and prevention of non-consensual image creation and sharing are necessary to keep pace with this growing threat to privacy, security, and democracy.

Increasingly, companies and governments are analyzing personal data to improve services, advance research, and combat discrimination. However, such analysis can also create valid concerns about differential treatment of individuals or harmful impacts on vulnerable communities. These concerns can be amplified when automated decision-making uses sensitive data (such as race, gender, or familial status), impacts protected classes, is used by courts or administrative agencies to craft sanctions or bestow benefits, or affects individuals' eligibility for credit, housing, employment, or other core services. Some technical, policy, and legal tools have emerged to identify and combat these risks.[6] But more progress must be made in identifying algorithmic risks - loss of opportunity, economic loss, social detriment, and loss of liberty – and potential mitigation strategies.[7]

c. **Data portability:** New data portability requirements are in effect in Europe (General Data Protection Regulation/GDPR), in California (California Consumer Privacy Act) and in many legislative proposals. Technical challenges include a lack of APIs that would allow any company to transfer data directly to any other company at the request of a consumer, and lack of any conventions for tagging/labeling data that would enable consumers or companies to make use of data they have "ported." Such technical challenges create barriers for companies complying with consumer data portability requests. Therefore companies rely on tools for providing consumer data that is structured, commonly used, and in a machine readable format.

Structured data is data from which software can extract specific elements. A common example of structured data is data contained in a spreadsheet in which specific elements can readily be accessed in an organized format.[8] Commonly used means a format that is widely-used and recognized across industry. However, not all commonly used formats are amenable to data portability as commonly used formats must also be structured and machine readable. Machine

---

[6] *E.g.* Sarah Tan et al., *Detecting Bias in Black-Box Models Using Transparent Model Distillation* (2017), *available at*: https://www.researchgate.net/publication/320464780_Detecting_Bias_in_Black-Box_Models_Using_Transparent_Model_Distillation.

[7] FPF, *Unfairness By Algorithm: Distilling the Harms of Automated Decision-Making* (Dec. 11, 2017), *available at*: https://fpf.org/2017/12/11/unfairness-by-algorithm-distilling-the-harms-of-automated-decision-making/

[8] ICO, *Right to Data Portability* (last visited Dec. 26, 11:17AM), available at: https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/individual-rights/right-to-data-portability/.

readable refers to data that can be automatically processed and read by a computer in a manner that is easily extracted or recognized.

While data that is structured, commonly used, and in a machine readable format is considered a strong prerequisite for effective data portability, other considerations aid companies in complying with user requests. For example, an interoperable format, which allows for data exchanged between companies to be readable by both companies, even if using different systems, would lead to efficient data portability.

Additionally, accurate verification and authentication, also known as entity resolution, of consumers making data portability requests is needed to not only comply with GDPR and California law, but also to ensure strong individual privacy for consumers from whom the data is derived. Verification and authentication occurs on two levels: 1. determining that the consumer making the request is the consumer they are claiming to be; and 2. determining that the data being requested is the data that is in fact associated with the consumer making the request. This two-pronged approach to verification and authentication may prove difficult for companies housing a trove of personal information, making it critical that companies access reliable software to complete data portability requests.

2. **Administrative/Compliance Tools**

a. **Risk Assessment Tools:** Both private and public entities face the challenge of prioritizing risks when performing functions and providing services using data. Risk assessment tools guide entities by: cataloging the types of risks that threaten privacy, utility, and trust; helping entities frame their primary objectives, while identifying the associated risks; prioritizing the associated risks; and suggesting appropriate controls for managing risk.[9]

Risk assessment tools should be implemented by organizations throughout the lifecycle of collecting, using, processing, and sharing data. This includes the earliest steps of framing organizational mission and objectives. Understanding the purpose of the business's data use is a necessary first step in determining how to respond to privacy and security incidents and selecting the most appropriate tools for safeguarding consumer data.

After framing organizational objectives, companies should identify the applicable obligations over consumers' personal data. Obligations occur on several levels including: legal and regulatory requirements; internal privacy policies or terms of service; applicable privacy-related principles, such as the Fair Information Privacy Practices (FIPPs); privacy-related goals as part of the organization's vision; and the organization's risk threshold and tolerance. When considering obligations companies should note the relevant types of privacy and security risks. While many privacy and security events cannot be anticipated by even the most punctilious engineers, categorizing and prioritizing risks will help internal engineering teams design systems to best

[9] NIST, *Privacy Risk Enhancement Methodology (PRAM)* (Mar. 28, 2019), *available at*: https://github.com/usnistgov/PrivacyEngCollabSpace/tree/master/tools/risk-assessment/NIST-Privacy-Risk-Assessment-Methodology-PRAM.

mitigate risks of the highest priority with the most vigor, while addressing lower priority risks as needed.

Throughout the process, organizations should assess internal system design with an eye toward privacy, while considering stakeholders' privacy expectations. Often referred to as "privacy by design," system engineers should access tools that take a holistic look at design throughout the entire lifecycle of a product or service, while ensuring that stakeholder and consumer privacy expectations are integrated at every step.[10]

b. **Data Sharing Toolkits:** Data sharing serves important social, economic, and democratic functions. However, shared datasets can carry risks to individual privacy. To encourage the socially beneficial use of shared datasets, while promoting strong individual privacy and addressing ethical concerns such as fairness and equity, data sharing toolkits provide guidance for navigating the complex policy, operational, technical, organizational, and ethical standards that support privacy-protective data sharing programs.[11] Toolkits can provide advanced statistical control strategies following a flexible, risk-based assessment process, and suggest standardized methodology to promote beneficial uses of data while addressing privacy concerns.

Although there is a growing body of research regarding open data privacy, open data managers and departmental data owners need to be able to employ a standardized methodology for assessing the privacy risks and benefits of particular datasets internally, without access to a bevy of expert statisticians, privacy lawyers, or philosophers. By optimizing its internal processes and procedures, developing and investing in advanced statistical disclosure control strategies, and following a flexible, risk-based assessment process, organizations can build mature open data programs that maximize the utility and openness of data while minimizing privacy risks to individuals and addressing concerns about ethical challenges, fairness, and equity.

Throughout the process of organizations adopting an open data program specific to the nature of the data, services, and community served, there are a number of privacy-related best practices in which organizations should engage, or at minimum consider. Organizations should document potential benefits and risks for each published dataset, both prospectively and retroactively, for those that have not yet had a benefit-risk assessment conducted. Organizations must develop policies and procedures for conducting additional screening of datasets and elevating the review of risky or sensitive datasets to disclosure control experts or a disclosure review board when appropriate. Finally, organizations must engage decision-makers at the data collection stage with decision makers at the data release stage (such as open data and public records staff), so that the full lifecycle of data collected by and for the organization can be better understood, managed, and communicated to the public.

---

[10] *See* Sophie Stalla-Bourdillon et al., *Data Protection by Process: How to Operationalize Data Protection by Design for Machine Learning* (Dec. 2019), available at: https://fpf.org/wp-content/uploads/2019/12/WhitePaper_DataProtectionByProcess.pdf.

[11] FPF, *City of Seattle: Open Risk Data Assessment* (Jan. 2018) at 3-4, *available at*: https://fpf.org/wp-content/uploads/2018/01/FPF-Open-Data-Risk-Assessment-for-City-of-Seattle.pdf.

Organizations should take additional measures as appropriate given the context of collection and sharing of open data, while utilizing toolkits to streamline the process, emphasizing privacy, utility, and affordability.

c. **Data Tagging:** Data tags are labels containing information applied to data. Legislation and legislative proposals are requiring companies to define uses and purposes related to data processing and to track data throughout a lifecycle and across partners. Moreover, when organizations enter data sharing agreements, contractual terms often require companies to accurately track data during exchange. Tools to support such data tagging requirements are nascent and will no doubt prove to be a necessity for organizations tracking data traveling through increasingly complex processing and sharing schema.

Not only is data tagging an important tool for complying with legislative and contractual requirements, it can also promote the public interest applications of datasets. For example, data tagging helps researchers replicate and authenticate scientific and medical findings. Additionally, companies benefit from data tagging to improve system processes and goods and services. However, some of the data shared in the public interest and for commercial purposes contains sensitive personal information. Data tagging can help organizations keep track of data to ensure appropriate disclosure, while avoiding privacy violations.

Data tags are individually tailored for each particular dataset to effectively label the data given its unique properties and obligations under the law or contract. Data tags are created and automatically applied to datasets by a system though asking an engineer a series of questions to understand the data's critical properties. The system then applies inference rules to the data to determine which laws, frameworks, or contracts are applicable to a particular dataset. Next, the system labels the data with specific information, such as icons, indicating the type of data, how the data should be shared, and other key properties.

While data tags are created for particular datasets subject to specific data sharing requirements, the informational icons generated through open source data tagging tools can be applied across industries. This standardization would create a holistic approach to data tagging, enhancing privacy and compliance across industries sharing data for multiple purposes.

d. **Entity resolution and authentication:** Entity resolution involves accuracy linking information stored on a database to its real-world counterpart. Requirements to provide access and deletion rights are driving demand for capabilities to collect data about an individual user that may be distributed across many sources, with incomplete information. Few tools to support these processes are available, making entity resolution and authentication, necessary for providing consumer privacy rights, difficult.

Tools for entity resolution should integrate five basic steps: 1. preprocessing data; 2. blocking; 3. matching; 4. verification; and 5. evaluation. Tools should verify with high-confidence that each step in the entity resolution process is met before proceeding to the next step. This can be

especially difficult because decisions must be made at each step in the process before moving to the subsequent step.

Preprocessing a dataset creates a standardized or common set of categories. Standardized data is created using elements such as, category order, punctuation, and character case norms. Additionally, each grouping of personal information is labeled with a unique I.D. to prevent mixing-up similar data elements.

The next step of entity resolution is blocking the data, which involves organizing pieces of personal information according to standardized fields such as “location,” which are unique enough to differentiate data elements. Data consisting of particularly sensitive data, or separate data elements that are similar, may be ripe for complex blocking techniques including advanced algorithms not readily accessible to most data scientists without the use of toolkits.

After preprocessing and blocking data, entity resolution tools engage in data matching, which matches data to the correct real-world entity. Data matching can be especially difficult for organizations attempting to match “messy” data, containing small differences in punctuation or spelling. Here, organizations must rely on increasingly complex matching tools to ensure accurate and efficient matching.

After data matching, organizations must verify that the matches made were accurate. Verification is an especially crucial step when dealing with a “messy” dataset in which matching all datasets is not possible. Here, organizations must engage in a benefit-risk analysis in deciding whether or not data matching is accurate enough to disclose the data when fulfilling a request. Finally, the evaluation step occurs in which a system, through a process of association, evaluates the broader landscape data relationships within a dataset to ensure that the overall system is correctly matching data. The need to correctly match data to a real-world entity is not only paramount to fulfilling a request, but also necessary to promote strong individual privacy and making certain data is only disclosed to the appropriate entity.

3. Self-regulatory and Policy Tools

a. **Industry best practices and codes of conduct:** Industry best practices, facilitated by an independent party, suggest practical solutions for maintaining individual privacy and promoting user trust in new technologies. Best practices provide guidance to companies grappling with the privacy and ethical concerns related to data use that are often left unaddressed by the law on the books. Additionally, best practices can serve as safe harbors under the law, reducing regulatory burdens, while preserving individual privacy without sacrificing innovation. Companies are often left with little guidance when developing or integrating emerging technology that cannot be mapped onto current law or legislative proposals.

Even if not covered by law, companies that prioritize strong privacy protections are often left with no guidance directly applicable to a particular technology or practice. Moreover, when guidance is available it often consists of a scattershot of principles and academic studies that sometimes

conflict with one another. Based on these considerations, companies are looking for strong best practices and codes of conduct applicable to a myriad of technologies.

Useful best practices can significantly vary in scope, from broad codes of conduct aimed at the technology landscape as a whole, to best practices providing guidance to a particular field, such as machine learning, to guidance for a particular technology, such as wearables[12] and facial recognition.[13] Regardless of the scope of a particular set of best practices, each framework must be flexible and adaptable to remain effective. Flexible best practices can be easily mapped to a host of technologies collecting, using, or sharing data in various ways. Adaptable best practices can maintain utility as a particular technology, or its associated concerns, evolves overtime.

Other than flexibility and adaptability, useful best practices should consider broad, fundamental privacy principles, such as the Fair Information Privacy Principles (FIPPS), regardless of how granular its provisions apply to a particular technology. Best practices will likely go beyond foundational privacy notions, but the inclusion or exclusion of foundational principles, or provisions going beyond foundational principles, should be a thoughtful undertaking. A wide-range of stakeholders should be involved in this drafting process to ensure best practices provide the guidance needed by stakeholders to achieve strong individual privacy.

Other than providing guidance, best practices and codes of conduct can serve as legally enforceable safe harbors, easing compliance burdens for companies and easing regulatory burdens for enforcement bodies. For example, companies looking to comply with the Children's Online Privacy Protection Act (COPPA) can rely on the Children's Advertising Review Unit's (CARU) code of conduct to comply with COPPA's children's advertising requirements.[14] Additionally, safe harbors often indicate to consumers that a product or service meets certain privacy-related requirements by providing an icon or seal indicating compliance.

b. **Independent ethical review boards:** One of the defining features of the data economy is that research is increasingly taking place outside of traditional academic settings. This includes companies, which may seek to advance societal causes or other agenda-driven projects. Independent ethical review boards operate as a standalone, on-demand review board to evaluate potential research uses of data and create a set of transparent policies and processes to be applied to such reviews. Independent review boards define the review structure, establish procedural guidelines, and articulate the substantive principles and requirements for governance. Independent review boards also address common company concerns about risk analysis,

---

[12] *E.g.* FPF, *Best Practices for Consumer Wearables & Wellness Apps & Devices* (Aug. 17, 2016), available at: https://fpf.org/wp-content/uploads/2016/08/FPF-Best-Practices-for-Wearables-and-Wellness-Apps-and-Devices-Final.pdf.

[13] *E.g.* FPF, *Privacy Principles for Facial Recognition Technology in Commercial Applications* (Sept. 2018), *available at*: https://fpf.org/wp-content/uploads/2019/03/Final-Privacy-Principles-Edits-1.pdf.

[14] *See also Student Privacy Pledge*, *available at*: https://studentprivacypledge.org/ (a public and legally enforceable statement by ed tech companies to safeguard student privacy).

disclosure of intellectual property and trade secrets when working with external researchers, and exposure to negative media and public reaction.[15]

Because of the dynamic nature of data science and its widespread impact, a diversity of voices and expertise are needed to tackle ex ante and ex post privacy issues. Independent ethical review boards composed of members with a diversity of expertise are well-situated to provide guidance. The range of expertise could include technologists, philosophers, lawyers, ethicists, marketing specialists, among others. Other than a diversity of expertise, independent review can provide a diversity of backgrounds and perspectives attuned to issues such as bias that are top-of-mind in emerging technology engaging in automated decision making, which are unlikely to be fully appreciated by internal teams alone.

While a recent industry push shows an interest in internal ethical review boards, outside expertise can provide guidance into how these internal ethical review boards are composed, how they function, how they define success, and when it is necessary to consult outside experts.

Finally, independent review boards can provide guidance and auditing services while avoiding conflicts of interest often at play with ethical review. Independent review boards that maintain confidentiality guarantees would not trigger the same intellectual property concerns companies encounter when sharing proprietary information with outside experts not part of an independent review board.[16] Confidentiality agreements with independent review boards can act as a gateway for a more fulsome analysis of a technology or service, leading to better informed ethical decisions by companies.

c. **Internal ethical review boards:** Businesses are designing, integrating, and implementing internal ethical review boards for emerging technologies, including automated systems, to build and maintain trust between businesses and customers through maximizing benefits and minimizing harms. Many of these ethical questions, such as fairness, bias, privacy, and economic impact, are not addressed by current law on the books because of a lack of regulation and inability to address ethical concerns in the black letter law. Internal review board frameworks provide initial questions companies need to be asking, such as who should serve on boards; how should boards be composed; guiding principles to inform ethical decisions; use of privacy enhancing technologies; and how to define success for a given product or service.

---

[15] FPF, *Conference Proceedings – Beyond IRBs Designing Ethical Review Processes for Big Data Research* (Jan. 5, 2017), *available at*: https://fpf.org/2017/01/05/conference-proceedings-beyond-irbs-designing-ethical-review-processes-big-data-research/ (summarizing an FPF discussion and workshop supported by NSF and the Alfred P. Sloan Foundation).

[16] Northeastern University Ethics Institute, *Building Data and AI Ethics Committees* (Aug. 2019) at 21, *available at*: https://cssh.northeastern.edu/informationethics/wp-content/uploads/sites/51/2019/08/AI-Data-Ethics-Committee-Report_V6.0-002.pdf.

Before an internal ethical review board can make determinations related to the development, release, and sale of emerging technology, a host of initial questions must be answered by the business. Development of toolkits containing key questions could act as a strong jumping-off point for a company adopting ethical review boards for the first time, or for new or evolving products or services. Initial questions should determine who will serve on the board, whether they be scientists, researchers, lawyers, or marketers, how many members should serve on the board, and where the board should be housed within the business.

After determining the make-up of the board, determinations surrounding the board's decision-making apparatus are needed to ensure decisions are made in a timely manner without ignoring input from various board members. This could include questions involving board voting, schedule of board meetings, feedback from outside experts, and various quorum requirements before making an important decision or revision.

After the board make-up and rules are in place, the board should begin with defining commercial, privacy, or ethical goals related to the product or service. This could include questions such as how the board will define success regarding both board process and meeting the goals set out by the board. Other questions the board should consider include: whether independent review is a necessary step; privacy by design principles throughout the lifecycle of a product or service; which, if any, privacy enhancing technologies should be integrated; and who is ultimately responsible if a technology goes awry, among many other considerations.

Because of the nascent nature of emerging technology, especially in the artificial intelligence and machine learning space, discussions surrounding internal ethical review boards are emerging at major technology companies. Many of these discussions pull from existing review boards and tools from the scientific and research community in an attempt to map successful frameworks that tackle difficult ethical questions to those posed by AI and machine learning.[17] Examples include, Institutional Review Boards (IRBs) required for human testing and frameworks used for animal testing. These established ethical review boards must examine ethical concerns such as consent, minimization of risk, and refining processes overtime. Internal ethical review boards should consider the applicability of existing ethical frameworks when developing emerging technology that poses ethical concerns.

Recommendations, next steps:

This report offers a snapshot from industry, produced by FPF as industry convener, of what technologies and tools are needed to safeguard personal privacy in an increasingly complex legal and regulatory environment. In addition to the privacy protective technologies described and business and research opportunities they represent, we make special note of the need to educate practitioners about research developments to build the field. The low level of awareness in the private sector of the academic state-of-the-art is referenced in the description of

[17] Sara R. Jordan, *Designing an Artificial Intelligence Research Review Committee* (Oct. 15, 2019) at 3-9, *available at*: https://fpf.org/wp-content/uploads/2019/10/DesigningAIResearchReviewCommittee.pdf.

de-identification tools where industry privacy leaders indicated significant interest in de-id but knew little about the current state of play. In a separate meeting with 15 companies that provide privacy tech tools for compliance, few of these companies were familiar with academic research or scientific advances in their particular focus area. As one of the academic observers at our New York City meeting commented, we need to address the academic to practitioner (and reverse) relationship "if we are serious about convergence."[18]

Private sector demand and readiness to collaborate is a factor in the launch and groundswell of activities associated with the Privacy Tech Alliance, an initiative established by FPF to advance privacy enhancing technology in commercial, government and not-for-profit sectors. The privacy technology track that FPF has proposed for NSF's Convergence Accelerator would encompass the Alliance and other related activities to purposefully integrate knowledge and expertise across multiple disciplines and sectors to balance requirements of data governance with individual privacy.

The track would seek to answer any number of business, research or policy questions about the future of privacy technology, possibly to include:

- Are specific emerging technologies best suited to specific industry sectors? Or geographical regions?
- How can privacy enhancing technologies and data protection by design methodologies be integrated into existing software development approaches, including especially agile development processes?
- What about users? How can we design and evaluate for users and for a democratic and equitable society?

Our proposal for a privacy technology track as the focus of NSF's latest Convergence Accelerator is driven by urgent organizational need across sectors to address the complexity and intensity of compliance and data governance while protecting individual privacy. Convergence effort of industry and academic researchers in this way would support the development and adoption of new technology with potential to fundamentally transform how data is used, managed and protected.[19]

[18] FPF NYC Peer-to-Peer Meeting (Oct. 24, 2019).
[19] *See* FPF, *Addressing a Critical Data Governance Problem through Privacy Protecting Technologies* (Jun. 18, 2019) (FPF's response to C-Accel RFI).

Appendix

Survey - The Future of Privacy Technology

In the face of shifting legal and regulatory frameworks, organizations in every sector of the economy are turning to privacy protecting technologies to balance the requirements of data governance and individual privacy. This survey asks industry privacy leaders to identify future directions of privacy technology to meet their business and policy needs. The results will inform the FPF's report to the National Science Foundation of business and research opportunities to fill gaps in the current privacy tech landscape and recommended action.

Question 1: What are the most significant privacy and data protection challenges that could be addressed by privacy enhancing technologies?

Question 2: What are the most promising privacy enhancing technologies that could mitigate privacy and data protection challenges? In the short term (1-2 years)? Medium term (3-5 years)? Long term (5+ years)?

Question 3: What are the most significant technical, organizational, or legal barriers to organizations adopting privacy enhancing technologies?

Question 4: Which privacy enhancing technologies are well suited to address challenges raised in particular industry sectors or with regard to certain data types?

Question 5: What are some examples of existing or emerging privacy enhancing technologies helping organizations comply with legal or ethical obligations? What's notably missing (where tech could fill a need)?

Question 6: With regard to de-identification technologies, what are the barriers to adoption or limits to the utility of these techniques today?

Question 7: Do you feel you are adequately aware of the leading scientific advances in PETs and whether they are feasible for implementation? Do you have methods to track such developments? What would the easiest way for you to stay informed and briefed about the state of developments?

About FPF

The Future of Privacy Forum (FPF) is a Washington, DC-based think tank that seeks to advance responsible data practices. The forum is led by internet privacy experts and includes an advisory board comprised of leading figures from industry, academia, law, and advocacy groups. For more information, visit www.fpf.org.