Internet Policy Research Initiative
Massachusetts Institute of Technology

Proposal: MIT Privacy Engineering Action Lab

October 5, 2023

Organization Information

1. Name of Organization

MIT Internet Policy Research Initiative (IPRI)

2. Discuss the founding and history of the organization.

Founding

The MIT Internet Policy Research Initiative (IPRI) is an MIT-wide initiative pioneering a new style of cross-disciplinary research and policy dialogue that brings together scholars from across campus. IPRI was created in 2015 with a founding grant of $15M from the Hewlett Foundation and early support from the Ford Foundation.

Since 2015, IPRI has established a track record of technology policy leadership and government engagement on issues such as encryption and surveillance, AI governance, cybersecurity, and privacy. IPRI's work on these issues has been strengthened by our trusted working relationships with technology policy leaders in the United States (both at a national and state level), the United Kingdom, Australia, Germany, and India, as well as with international organizations, including the OECD and the World Bank.

Daniel J. Weitzner, holder of the 3Com Founders Senior Scientist chair at at the MIT Computer Science and Artificial Intelligence Lab (CSAIL), founded IPRI in 2015 as a response to the critical need for technology-informed policy making in the areas of privacy, security, networks and the Internet economy. The group plays an important bilingual role of informing policy making with technical expertise, and helping engineers build secure and privacy protecting products that are informed by policy.

IPRI is in a unique position to advance individual privacy rights through computer science research that will create new privacy-preserving technologies, and public policy research to explore technically-grounded advances in privacy policy and law. IPRI's senior leadership has strong consumer and Internet civil liberty advocacy backgrounds. Daniel Weitzner was the first staff member in Washington DC for the Electronic Frontier Foundation and founder of the Center for Democracy and Technology. He was also a senior policymaker (White House Deputy CTO for Internet Policy). While at the White House, Weitzner was responsible for developing

the Consumer Privacy Bill of RIghts in 2012. Dr. Taylor Reynolds, IPRI's Research Director, was the senior economist at the OECD responsible for the Internet economy, and his research on broadband pricing led to multimillion dollar fines against incumbent telecommunication firms engaged in deceptive advertising.

Of particular relevance, Daniel Weitzner has a long history of successful Internet civil liberties advocacy. His work led directly to amendments to the Electronic Communications Privacy Act in 1994 that offered groundbreaking protections for web browsing logs, email records, and other transactional data. (18 USC 2703(d)) Under Weitzner's leadership, the interests of the class in better privacy protection will be materially advanced.

Our research streams

IPRI's primary research efforts have historically covered six core research areas: cybersecurity, privacy, networks, AI policy, the Internet experience, and AppInventor. The IPRI team is led by a core group of 16 principal investigators. Over 20 students are involved in IPRI research at any given time. Our structure as an cross-campus entity that brings together technical experts from across fields with policymakers to tackle key challenges has produced significant societal impacts in health, voting, and privacy among others. Several of these are highlighted below:

Private Automated Contact Tracing (PACT)

- **Technology/Policy Problem:** There was a need for contact tracing applications during COVID-19 that were secure and preserved privacy.

- **IPRI Solution:** IPRI put together a consortium called PACT (pact.mit.edu) that brought together technologists and public health officials to automate parts of the exposure detection function while maintaining user privacy and ensuring equitable deployment.

- **Impact:** PACT-designed exposure notification technology, including state-of-the-art security and privacy architecture, is now included in billions of Android and iPhone products all around the world.

Voting Apps

- **Technology/Policy Problem:** There is an ever-increasing interest in online and mobile voting, but the cybersecurity profile of such services raise numerous questions.

- **IPRI Solution:** IPRI researchers published a technical paper demonstrating deep security vulnerabilities in a particular mobile voting app (Voatz), the first thorough reverse engineering analysis of a live mobile voting system to demonstrate many of the security risks long warned-of by the computer security community.

- **Impact:** Several states and other voting jurisdictions that either were using Voatz or considered doing so reversed their decisions after IPRI researchers shared their findings with the Department of Homeland Security's Cybersecurity and Infrastructure Agency (CISA), Voatz themselves, and later major news publications like the New York Times.

Encryption and surveillance

- **Technology/Policy Problem:** Law enforcement agencies around the world propose redesigning Internet and smartphone infrastructure to enable government access to encrypted information. But is this secure?

- **IPRI Solution:** The Keys Under Doormats Report, produced by the world's leading cryptographers and cyber security experts showed that enabling "backdoor access" posed security risks.

- **Impact:** The report findings changed the direction of U.S. policymaking, and encouraged Australian and UK parliaments to update their legislation to address systemic risk identified by the IPRI report. IPRI remains engaged in research and dialogue with government agencies, companies and civil society to identify constructive approaches with reasonable risk levels.

Lack of cybersecurity data for decision and policy making

- **Technology/Policy Problem:** Little data exists about the most efficient means of addressing cybersecurity risk, so investment decisions are suboptimal, cyber insurance is inefficiently priced, and policymakers lack the information necessary to make sound public policy. This lack of price/risk data is because firms don't share data on cyberattacks, so we learn nothing about the attacks or the effectiveness of cyber defenses.

- **IPRI Solution:** A team at IPRI developed the SCRAM (Secure Cyber Risk Aggregation and Measurement) benchmarking platform that uses secure computation techniques (multi-party computation) to securely and privately aggregate sensitive data without requiring disclosure of the underlying inputs.

- **Impact:** The IPRI team produces new cybersecurity benchmarks, metrics, and models that are used to forecast cyber risk for various sectors of the economy. These inputs are then used to create policy recommendations and guidance for municipalities that operate critical infrastructure and federal policy making bodies.

AppInventor

- **Technology/Policy Problem:** Mobile phones have become an important interface for sharing and consuming information, but the process historically for creating apps for mobile platforms was difficult and expensive, particularly for socially beneficial applications with little to no available funding.

- **IPRI Solution:** The AppInventor team built an intuitive, visual programming environment that allows everyone, even children, to build fully functional apps for smartphones and tablets. Those new to MIT App Inventor can have a simple first app up and running in less than 30 minutes.

- **Impact:** AppIventor has over 400,000 unique monthly active users who come from 195 countries who have created almost 22 million apps. MIT App Inventor is changing the way the world creates apps and the way that kids learn about computing.

3. Describe the organization's current goals.

The Internet Policy Research Initiative's (IPRI's) mission is to lead the development of policy-aware, technically grounded research that enables policymakers and engineers to increase the trustworthiness of interconnected digital systems like the Internet and related technologies.

To achieve this mission, IPRI produces fundamental, cross-disciplinary technology and policy research (publishing 35 research papers in 2022); engages with global policymakers, industrial partners, and civil society organizations; and is building a network of students educated in the field of Internet policy.

4. Provide a brief description of the organization's current programs.

MIT is one of the top universities in the world across a number of disciplines, including engineering, computer science, and economics. MIT has 11,376 students and 13,000 employees. Recently the Institute announced the creation of the Schwarzman College of Computing which represents a new paradigm for computer science research and education that recognizes the importance of addressing the social, ethical and policy impact of computing on society. Currently, IPRI has six main research streams.

<u>IPRI by the numbers in 2023</u>

- PIs: 16
- Students doing research: 41
    - 7 PhD
    - 15 masters
    - 21 undergraduate

- Publications: Roughly 35 per year

Current research streams

1. ***Privacy,*** covering topics such as designing new databases and systems embedded with privacy protection and user control, evaluating the international privacy policy landscape and studying privacy incentives, data protection policy, web surveillance, human-computer interaction in the context of privacy, the implications of silently listening, and overarching insight into the global privacy research area.
2. ***Cybersecurity,*** covering topics like encryption policy, accountability, cryptography, data sharing, securing core economic and social infrastructure, and measuring cyber risk.
3. ***AI Policy,*** covering topics like the role of AI in financial decision-making, increasing access to new training data sets with policy, working with stakeholders on AI principles, and shaping global Internet policymaking via policymaker engagement and informing the public debate.
4. ***Networks,*** covering topics like Internet architecture, Internet security, Internet economics, Internet policy, and network management.
5. ***Internet Experience,*** covering topics decentralized privacy preserving platforms for clinical research, the trustworthiness of autonomous systems, the relationship between privacy and machine learning, complex machine and model explanations, securely aggregating distributed data, and developing smart contracts for data sharing.
6. ***App Inventor,*** involving the creation of a tool to enable anyone, especially youth, to develop mobile apps that better their communities.

In addition to the high-impact research activities described above, other relevant IPRI contributions include:

- We developed “Privacy Bridges” with European partner universities to help create a framework for data protection and usage between the US and the EU. Our report was presented at the International Conference of Privacy and Data Protection Commissioners.
- Our team contributed technical and policy guidance to the OECD as they developed the OECD’s AI Policy Principles that were adopted by 36 countries. IPRI sent three experts to participate in the OECD’s Expert Group on AI. IPRI also hosted the OECD’s AI Expert Group Meeting in January 2019.
- Daniel Weitzner was selected to be a member of the OECD’s Expert Group to revise the OECD’s long-standing privacy guidelines.
- Our researchers and leadership frequently prepare submissions to governments related to encryption policy. Our researchers have testified in front of the US Congress and were invited to testify before the Australian Parliament on these issues.

5. **Has your organization ever received a prior cy pres award? If yes, please cite the applicable case(s), identify the amount(s) awarded, and describe the nature and scope of the project(s) funded.**

IPRI has received a cy pres award in the matter of In re Google Inc. Street View Electronic Communications Litigation, No. 10-md-2184, Northern District of California) in the amount of $1,006,582.

6. **Has your organization been reviewed or rated by Charity Navigator or similar entity?**

Yes. MIT has a Charity Navigator rating of 96%.
https://www.charitynavigator.org/index.cfm?bay=search.profile&ein=042103594

7. **Identify Principal Investigator/Project Director**

Daniel J. Weitzner is the MIT IPRI Founding Director and holds the 3Com Founders Senior Research Scientist chair at MIT Computer Science and Artificial Intelligence Laboratory.

8. **Provide a summary of the plan for the program or project request. Include the issue and/or opportunity addressed, goals and objectives, activities, and timeline.**

Introduction and Motivation:

The MIT Internet Policy Research Initiative proposes to launch a new MIT Privacy Engineering Action Lab (PEAL). PEAL will materially advance the interests of the class in In re Google Location History Litigation, No. 5:18-cv-05062-EJD (N.D. Cal.), helping to assure than members of the class, and those similarly situated in the future are far less likely to be victims of privacy harm arising from deceptive collection of personal data, the inability of users to control how their data is used, and the general lack of technical tools and design patterns that encourage respectful privacy practices.

As serious as are the harms in this case, they are only the tip of the iceberg. Internet users are surrounded by systems with equally deceptive, uncontrolled and virtually-invisible flows of personal data in the fields of health, finance, transportation, employment and other sensitive activities of our daily lives. What's more, these privacy risks are growing over time: personal data flows are becoming more complex, with many organizations moving to business models in which sensitive personal data is passed rapidly across organizational boundaries and subject to increasingly penetrating analysis through a growing variety of statistical and AI machine learning techniques. As these privacy risks grow, users' ability to exercise control over their data, and regulators' ability to detect wrongdoing are all declining, putting users at even greater risk.

We propose a coordinated package of computer science research, user education and professional development materials that will provide new technical accountability measures, quantitatively rigorous user studies that identify both new dark patterns and recommended affirmative design practices, and new educational approaches for computer science undergraduates as well as professional development education for a developing new job category, Privacy Engineers, who will contribute to much of the technical privacy environment for users in the future.

The project is divided into three components, each running concurrently for three years.

1) **Traceable Accountable Privacy Protocol (TAPP)**: Robust technical infrastructure that processes personal information at web scale in a demonstrably reliable manner. This part of the project will engage computer science faculty and students to develop a new web protocol that will facilitate sharing personal data in a fully-accountable manner, so that each user is able to attach clear consent and usage conditions for their data, know who has their data, what they are doing with it, and easily cause it to be deleted or corrected. We will publish a working description of this new protocol and make available open source libraries to enable its widespread implementation and adoption. Members of the IPRI team have extensive experience moving technical designs from the lab to the web, including PI Weitzner's leadership role at the World Wide Web Consortium (the body that sets technical standards for the Web) as well as Weitzner and Liccardi's experience developing the COVID Exposure Notification Protocol (PACT) ultimately implemented on billions of Android and Apple iOS devices during the pandemic.
2) **Learning From User's Behavior 'in-the-wild'**: We have endless privacy surveys, all of which establish that users want more control over their data and distrust many who have it today. By studying how users actually interact with different services, we will identify dark patterns in new styles of user interfaces such as voice interfaces and AI-powered chat bots. We can also develop positive user interface design patterns that encourage respectful handling of personal data. This component of the project will produce authoritative behavioral science studies (using Human Computer Interaction techniques) to identify with precision design features that are deceptive, and use these scientific insights to inform design recommendations which will be made widely available to the developer community.
3) **Privacy Engineering education**: Enterprises large and small, public and private, all face the challenge of how to handle personal data in a respectful and accountable manner. Ultimately, the technical side of this challenge will be met by a new category of engineers and software developers – the Privacy Engineer. We propose to create educational materials to guide the development of this new field, both for use in undergraduate computer science education and for those already operating in professional settings. Of course, the decision and incentive to be responsible about privacy begins with the legal system and a commitment to ethical business practices.

Once firms make this commitment, they will turn to their privacy engineers to do that work, so we want to contribute to their being ready for the challenge. MIT IPRI has been contributing to this field and studying its needs[1], so are well-positioned to contribute to its growth.

Individual Project Descriptions

Project 1: Traceable Accountable Privacy Protocol (TAPP): A new web protocol for sharing personal data in a traceable, accountable fashion

As companies collect and share vast amounts of consumer data, trust in their data stewardship practices is rapidly declining.[2] This trend is especially worrying given that in many sectors there are moves to share more, not less, personal data, and that data will increasingly flow across organizational boundaries. This will only increase consumer confusion and decrease individuals' practical ability to control their data. What's more, key sectors where sharing is increasing – healthcare, finance, employment – pose a risk for harm as great or greater than what we have seen in social media and online advertising. While cross-enterprise data sharing can bring consumer benefits and convenience, it has also led to many high-profile consumer data protection violations, exemplified recently by the Facebook-Cambridge Analytica scandal. This was among the largest known bulk misuse of personal data and it happened precisely because Facebook was careless about allowing data to flow off its platform without any controls. What's more, that fact that Facebook was under an FTC consent decree with bi-annual audit provisions raises doubts about how well our enforcement processes are working[3]. As we move toward a world in which data sharing is only likely to increase, we need better technical tools to ensure user control and accountability, along with better legal protection and enforcement. At MIT we have limited ability to influence the state of privacy law and enforcement, but do have the ability to design and deploy new privacy protection technology.

In order to regain consumer trust and provide meaningful control, it is essential for organizations to:

- obtain explicit, informed, and granular consent when processing personal data, which will give consumers greater clarity over the type of information they share and how it is used;
- offer traceability in their data use and sharing practices, which will give users control, as they know who has access to their data and how it is actually being used;

[1] Gulati, Liccardi, Weitzner, "Privacy Law in Practice: Exploring Challenges to Modern Privacy Compliance," Privacy Law Scholars Conference, June 2023.
[2] Kimberly Bella, Christophe Carugati, Cathy Mulligan, and Marta Piekarska-Geater. Data for common purpose: Leveraging consent to build trust. https://www3.weforum.org/docs/WEF_Data_for_Common_Purpose_Leveraging_Consent_to_Build_Trust_2021.pdf, 2021
[3] Weitzner, How Cambridge Analytica, Facebook and Other Privacy Abuses Could Have Been Prevented, Lawfare, April 4, 2018.

- provide mechanisms for accountability in case of any violations, which will give consumers confidence that those holding their data will be held accountable for any misuse and that regulators can identify abuse of personal data.

To give consumers a more complete understanding of how companies handle their data, we will design and prototype a suite of data governance protocols to facilitate consent, traceability, accountability, and portability across enterprise boundaries. This will unify how consumers manage consent and data lifecycles with all of their service providers, while also standardizing how different providers report on data use and sharing events. Our design goals are focused on simplicity (they should rely on established standards where possible), security (they should adopt best practices for secure data storage and transmission), and scalability (they should allow companies to scale out as their userbases and databases grow).

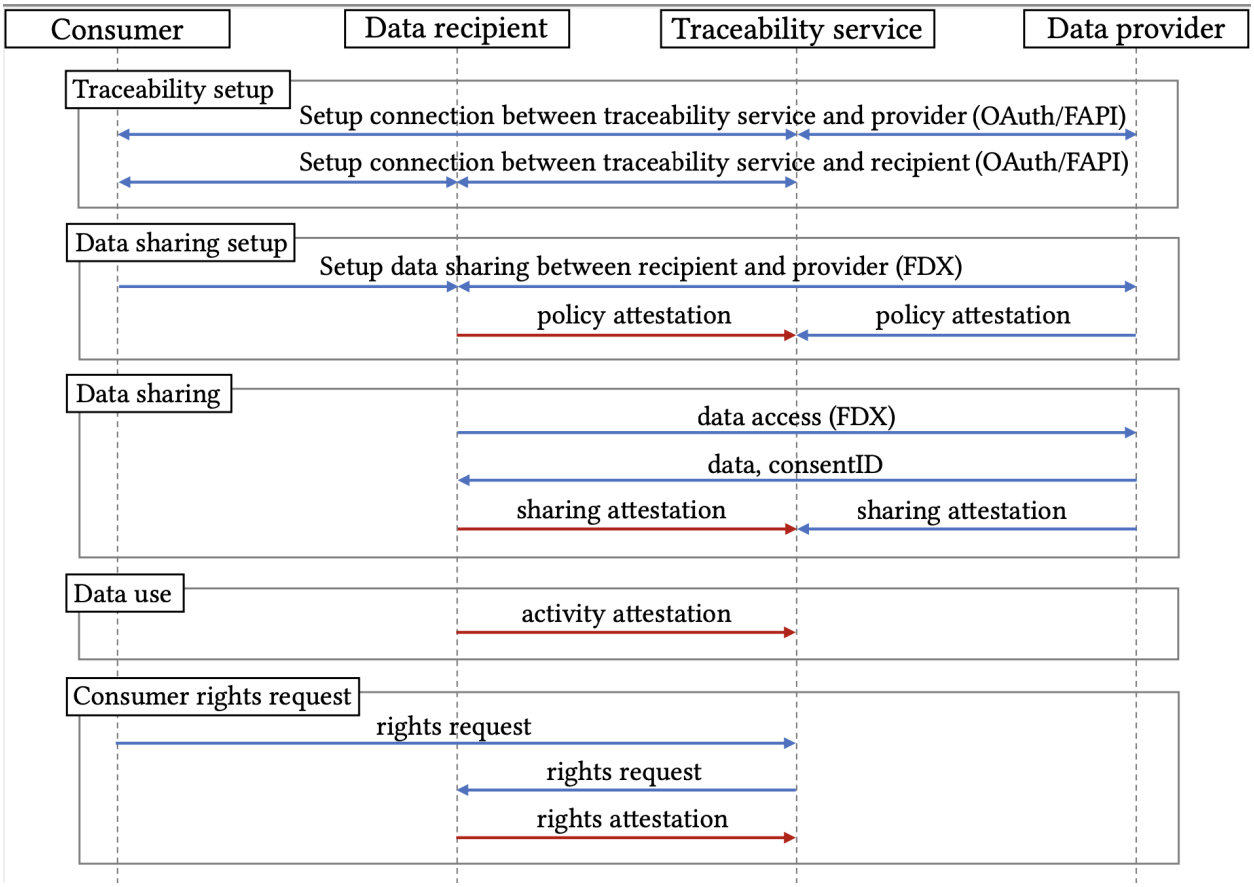

The figure above is our initial design sketch of this technical protocol. The key idea of this protocol is that each time data is shared with a third party, a clear statement of consumer consent should follow with it, and every time a third party uses that personal data, a record of that use should be provided to the consumer or her agent, with the ability for the consumer, her agent, or even a regulator, to easily (with machine assistance) assess whether the uses are permitted. The same mechanism should give consumers the ability to know who has their data

at any moment, as well as a mechanism to have the data deleted or corrected upon the instruction of the consumer.

Opportunity: There is a unique opportunity to change the terms on which personal data is handled because a number of business sectors are making major technology investments in new data sharing technologies. Now is the time to add explicit privacy protection and user control features while the new services are still in the design phase. For example, many banks are investing in new open banking services which will entail larger-scale and more complex flows of personal financial data. Regulators around the world are also considering modernizing personal data rules. It is essential to show that user control and accountability can be added along with these new services in a manner that enhances user privacy rights, improves the possibility of meaningful privacy enforcement, and still enables consumers to have the benefit of these new services. Our experience in design and deploying web standards and new privacy protocols in other settings positions IPRI well to have a positive impact on the privacy technology landscape.

Funding: Over the course of a three year project, we propose to fund 3 graduate students, 2 undergraduates, one postdoctoral researcher, a small part of one senior faculty member in computer science, all supervised by PI Daniel Weitzner. The cost of this component of the project is $899,000 over a three year period.

Target population: The target population for this component of the project are end users of consumer-facing services on the web such as banking, healthcare and other complex data analytic platforms. The ultimate impact of this project, if successful, would be to ensure that all personal data flowing through these new services is under the control of individual users and handled in traceable, accountable fashion.

Project 2: Learning From Users' Behavior In-the-Wild: User Experience studies

While we know users want control and we know that trust emerges from better control and more transparency, the details of what kind of control and how much information to offer users is not always well understood. The best path to designs that have these properties entails observing (under Institutional Review Board ethics control) users as they use real services to understand their behaviors and motivations. Using Dr. Ilaria Liccardi's experience sampling studies in the wild over a period of time, we will develop clear design patterns that meet user privacy expectations and build trust. These studies will begin with design sketches presented to focus groups, develop into pilot projects that observe users interacting with real world services, and ultimately mature into larger-scale deployments that engage with users as they are using new privacy features in live systems.

Looking forward, we must also extend our understanding of privacy dark patterns and respectful interface designs into other user interaction methods. Voice-activated interfaces and AI-powered chat bots are the new mode through which many users will interact with next-generation digital services. Given the intrusive data collection capability of these interfaces, we must understand the privacy risks and how to mitigate them.

The widespread adoption of smart assistants, both within the home and on smartphones, has made audio open to possible intrusion by technology providers. We propose to investigate the privacy implications of audio captured around voice activate assistants on smartphones and on stand alone devices. We aim to understand the types of audio that can be captured – across different locations – and to explore whether people understand the types of information that they are giving away when using these devices. We aim to compare people sharing preferences before and after exposing the implications of using these devices to people's privacy. In our previous work we found that regulations are lacking and in need to be amended to include possible privacy violations of these devices.[4] In particular, given the ubiquity of these devices, regulations should be considered, especially because our previous research[5] showed that these devices placed in the home capture less intrusive and privacy sensitive information compared to portable devices such as apps on smartphones and/or car applications.

This work can also be extended and provide insights into the privacy implications of AI-powered chatbots given their increasing widespreads use. Similarly to audio based commands, the privacy implications of using AI-powered chatbot applications are often opaque, unclear and misunderstood by users. We aim to investigate users' usage over time after possible implications are highlighted and communicated clearly. We believe that users' behavior and usage of this tool might change when possible intrusions are presented.

The result of this project can highlight people's actual preferences when it comes to their personal data. It can guide companies to create, adapt or re-evaluate how to use the information captured. It can help regulators create and amend regulations to safeguard people's data around these devices.

From this project we will produce several peer-reviewed papers reporting on the results of our user studies. These papers will both identify new styles of dark patterns and also recommend positive best practices for respectful privacy interface and interaction design. This rigorous scientific evidence will also help guide privacy enforcement authorities and others in the legal system to hold irresponsible data controllers to account.

[4] Lindsey Barrett & Ilaria Liccardi, Accidental Wiretaps: The Implications of False Positives by Always-Listening Devices for Privacy Law & Policy, 74 Okla. L. Rev. 79 (2022)
[5] Ilaria Liccardi & Jose Juan Dominguez Veiga, Wiretapping Your Friends: Privacy Implications of Voice Activated Assistants (on file with authors).

Funding: Over the course of a three year project, we propose to fund 2 graduate students, 3 undergraduates, one postdoctoral researcher, a small part of one senior faculty member in computer science, all supervised by Dr. Ilaria Liccardi and PI Daniel Weitzner. We will also have to fund the cost of working with a large number of users as study subjects. Best practice is to compensate those who are willing to participate in these long-run user studies The cost of this component of the project is $937,000 over a three year period.

Target population: The target population for this component of the project is end users of consumer-facing services on the web, especially those that use new voice-activated and AI-backed user interfaces. We also aim to influence the designers of these services and provide guidance to those who regulate them.

Project 3: Privacy Engineering Education

The long run privacy welfare of those who use digital services depends on the privacy norms that society sets and enforces through law and community demands. Designing and implementing the ever-evolving new digital services will depend on the engineering talents of software developers in general and on a new class of engineers known as Privacy Engineers. We propose a two-pronged educational approach that will enhance the privacy awareness of computer science students generally, and to help inform the development of the technical and professional standards of the new sub-field of Privacy Engineering.

Privacy awareness for computer science students: We will build on the novel, multi-disciplinary education approach of MIT's Internet Policy Research Initiative by extending two courses currently offered by IPRI faculty: 6.4590: Foundations of Internet Policy, and 6.S978: Privacy Legislation Law and Technology (offered jointly between MIT Electrical Engineering and Computer Science Department and Georgetown Law School (see New York TImes: Natasha Singer, Top Universities Join to Push 'Public Interest Technology', March 11, 2019; MIT Spectrum, Legal/Code-MIT engineering students team up with Georgetown lawyers-in-training on internet privacy legislation, Winter 2018). These courses teach 30+ computer science and engineering students each semester to develop the intellectual skills necessary to understand the complex public policy questions, including privacy, raised by computing in our society today. PEAL will add a hands-on laboratory component to each course giving students an in-depth experience of actually building and analyzing technical systems that address privacy harms.

By expanding these well-established courses, we will give our students added engineering experience needed to design and develop applications using personal data in a manner that does a better job of adhering to privacy law and best practice, thereby avoiding privacy harms suffered by the class of plaintiffs in this case. Engineering students learn good software development style through practice. We already have a well-developed curriculum for teaching our students how to understand broader issues of law and policy. By adding lab components to

the courses, we will give students concrete software development challenges that test their policy knowledge and give them the experience to make good design decisions in their careers. To help students understand and master the challenges of privacy-aware system design, we will build software platforms that simulate large-scale databases of personal information as environments within which students can experiment with different privacy designs. Developing lab teaching materials is a resource-intensive task, so support from this fund will be critical. IPRI will hire additional teaching assistants and a postdoctoral fellow to supervise the development of the new lab materials. Once these are developed, however, we will make them available freely to the rest of the academic community and professional software developers around the world.

Engagement with Privacy Engineering discipline: We know from studies of the development of the privacy engineering field that individuals who take on these jobs are passionate about building privacy-respectful systems. We will develop a series of monthly professional development seminars for early- and mid-career privacy engineers, providing them background on the latest research from our lab as well as our colleagues in universities around the world. We will also create a forum for peer interaction among privacy engineers. We know that there have been efforts to create such fora on a commercial basis that have failed. We believe MIT can be valued, trusted convenor for this new discipline.

Funding: Over the course of a three year project, we propose to fund 1 graduate teaching assistant to develop course materials, as well as supplementary faculty and lecturer time to create a new syllabus. We will also have expenses in developing and distributing the remote interactive course materials. The cost of this component of the project is $631,500 over a three year period.

Target population: The target population for this component of the project is computer science undergraduates at MIT and other universities around the world, as well as privacy engineers.

9. Explain why the organization is approaching the issue and/or opportunity in this way.

[See individual project descriptions for opportunity assessment and discussion of organization approach]

10. Identify and explain the range of funds required to effectuate the program or project request, on an aggregate and annual basis (if applicable), including how the money will be used.

We request a total of $2,467,500 to be expended over approximately a three year period. Each individual project described in paragraph 9 above shows specific funding amounts and spending categories. We summarize here:

PEAL Project Component	Funding (over 3 years)
1: Traceable Accountable Privacy Protocol (TAPP)	$899,000
2: Learning From Users' Behavior In-the-Wild	$937,000
3: Privacy Engineering Education	$631,500
Total	$2,467,500

11. Will the money be used to continue an existing project or create a new project?

PEAL will be a new activity that is part of the MIT Internet Policy Research Initiative (IPRI). IPRI is funded by the William and Flora Hewlett Foundation Cyber Program with a leadership grant of $15 Million.

12. What target population will your organization's project benefit?

[See individual project descriptions for target population of each component of the project.]

Evaluation

13. Will your organization agree to provide a report to the Court and the parties every six months informing the Court and the parties of how any portion of the Settlement Fund allocated to it has been used and how remaining funds will be used?

Yes

14. Describe how your organization will evaluate the success of the grant on enhancing or promoting the protection of internet privacy.

Overall, our most important measures of success are three-fold: First, we aim for widespread deployment of the new technical privacy protocols we are designing. If these protocols are used in a single large sector of the Internet such as financial services, healthcare or social media, we will consider our work to have a high level of impact. Even if this is not the case, we may see aspects of our protocol design in widely deployed software and services. That would also be a success. As an interim measure of impact, we will measure the number of developers who download the libraries we put out. Second, we will consider our work on user interface analysis and design to be successful if (a) our identification of new dark patterns are used in public policy development, government enforcement actions, or private litigation on behalf of individuals suffering privacy harm from deceptive practices. Finally, we will measure the impact of our

education programs based on the number of privacy engineers who use our educational materials. As soon as those materials are developed, we will provide usage statistics in our semi-annual reports.

15. Does your organization intend to use the results of the project in any publications, conference papers, and presentations?

Our research results will be published in leading peer-reviewed academic journals. Once research has been peer reviewed for quality, we will also seek to publish short-form versions of our work general audience publications such as Lawfare, the Harvard Business Review and op-ed pages so that our work receives broader impact. The technologies we develop will all be made available under open source licenses to encourage widespread usage.

* * * * *