Centre for Information Policy Leadership (CIPL) and IEEE Digital Privacy Initiative

Privacy Engineering: Aligning Technology, Principles, and Governance Roundtable

Event Takeaways | November 2025

Centre for Information Policy Leadership (CIPL) and
IEEE Digital Privacy Initiative

Privacy Engineering: Aligning Technology, Principles, and Governance Roundtable

Event Takeaways

As organizations grapple with increasing consumer privacy expectations and complex regulations, the need for collaboration between privacy engineers, product designers, marketing departments, legal teams, and policymakers has never been greater. In response, the IEEE Digital Privacy Initiative and Centre for Information Policy Leadership (CIPL) hosted a forum that brought together key stakeholders across sectors to explore organizational aspirations and operational challenges in developing products and services that respect privacy while enabling responsible and beneficial use of data. To encourage open discussion, the roundtable was conducted under the Chatham House Rule. Following are initial takeaways from this interactive conversation, organized by discussion segments.

Opening Keynote on Data Strategy

Data has never been more important for the success of businesses and organizations across every sector, and **a company's data strategy should be integral to its business strategy**. Success can be driven by realizing the potential value of data while effectively managing risk; if your data strategy is built to simply meet basic compliance, it is missing significant opportunities to benefit from responsible data use.

Collecting, processing, and analyzing data with clarity on how it can support business objectives, in addition to taking steps across the data lifecycle to ensure data quality and security, helps to limit risk while increasing the potential for value. Strive for a 70/30 value-to-risk ratio and invite regulators into the conversation. Try to look ahead and plan for future regulatory requirements rather than focusing solely on the current status quo.

Also consider customer trust metrics; it is important to center the customer in any business strategy, understanding that how personal data is processed is the basis for establishing trust. Reference the IEEE Code of Ethics for engineers "to hold paramount, the safety, health, and welfare of the public." CIPL's Accountability Framework is another useful point of reference. Focusing on value, including consumer privacy, can act as a key market differentiator.

Challenges identified for future discussion:

1. How do we measure risk accurately and usefully?

2. How might organizations select metrics that not only deliver and demonstrate tangible value, but also align with organizational values and an ethos of organizational accountability?

Discussion on "Privacy Expectations: Getting Stakeholders to Work Together"

Ensuring that all relevant teams are involved in the process from the outset, rather than after a design is finalized, is crucial. One of the most challenging tensions that privacy engineers must navigate is that different teams (e.g., marketing vs. data security) will have different priorities when it comes to the amount and types of data to collect, as well as the extent of the measures to protect that data.

Making the case that better privacy measures enhance consumer trust may strengthen the argument for privacy engineering approaches. The business focus is often on speed and innovation, and some decision makers may have a higher tolerance for risk as a tradeoff for the speed of innovation, even if it bypasses privacy. However, customer perception with regard to privacy and security may be a competitive advantage.

Data is meaningless unless you understand the context and can derive value from it. Instead, the focus should be collecting and processing the minimum set of data that is truly needed for the task at hand. Privacy engineers can help to operationalize innovative approaches to data minimization, where sufficient data is collected to achieve business purposes while high privacy standards are upheld.

Consider starting from practical use cases, which will help highlight the needs/requirements of privacy tools in tangible contexts and circumstances. When privacy discussions are theoretical, the conversation on privacy tools becomes stale very quickly. The moment privacy impacts business processes, things change.

Take a multi-dimensional approach to assessing the customer experience. User experience can often take a backseat to engineering, but putting the customer in the center of the design process is essential. Talk to users of the app/tool but also consider looking at their experience from a different perspective: Who are the non-users of your app/tool? What about your privacy measures may have made them choose not to use the app/tool?

Getting people on the same page requires having a page. Decide together on your privacy guarantees and guiding principles, and make sure these are aligned and understood across the company. Refer to this document in discussions. Ask for features to be reviewed, as you likely know what would be approved under these guidelines. Post your company's statement for public knowledge and awareness.

Consider ways to develop privacy tools that simplify the process and give individual users more control. The Notice and Consent approach is ineffective and often imposes counterproductive friction. Users want to download an app without the need to read long legal documents or consult a lawyer. The focus should be on building consumer apps with privacy and individual experience in mind, featuring user-beneficial designs that limit scrolling and utilize visual elements.

Frame digital product safety for consumers in the same way as physical safety. Consumers may have difficulty contextualizing risks associated with the use of their data, and they may have limited time or expertise to evaluate inherent privacy risks. Make these risks tangible for consumers, for example, by clearly conveying that identity fraud may cause significant financial harm. Recent studies have shown that once privacy-protecting mechanisms are enacted, fraud is reduced. This is one way to illustrate the consequences and benefits to both decision makers and users.

Consider best practices for building a compliant product. Examine the legal landscape, encompassing state and federal laws, as well as international statutes and agreements. These combined are the foundational standards that a product will need to be compliant. To meet these goals, look at a product that already works within these constructs and consider internal incentives for the engineers to work to meet these standards. You may also need to identify where laws and regulations are in contradiction and develop a risk-based approach to manage these issues.

Acknowledge that compliance with the law is the bare minimum and that laws do not necessarily reflect what the people want. Having a set of principles that the company stands behind can build trust with users. These can also serve to enhance business motivation and brand differentiation, moving beyond legal compliance.

Examine possibilities for building certified solutions. Software development is generally not regulated with certification and licensing. This underlines the need for a set of guiding corporate principles, such as a Code of Ethics or similar guidelines. Privacy engineers need standards to benchmark their work against established norms. (See also IEEE Standards; for example, IEEE P7002-2002, which integrates privacy into Software Development Life Cycle [SDLC] processes.) Consider how best to communicate to users when the company is certified and doing what it's supposed to, especially in light of growing concerns surrounding AI.

Define your North Star. Develop a strategy for the organization to make it easier not to get lost in implementation. Be more proactive and less reactive. Map the gaps and risks.

Discussion on "From Principles to Implementation: How to build products that respect privacy with a focus on AI"

Embed privacy into the DNA of AI. Begin by having privacy experts at the table during the design phase rather than bringing them in once a model exists. Consider strategies for putting privacy into all these systems and minimizing data from the beginning, by labeling early on and enabling purpose-based training. Measure the risks as well as the effectiveness of guardrails.

Regulators and regulation can support privacy by design and sound AI governance practices. Customers want measures that enable privacy by design and sound AI governance, and regulators can support such measures through incentives, guidance, and mechanisms for responsible experimentation and innovation, such as regulatory sandboxes.

Technical approaches to managing the appropriate use of personal data in model training are still in their infancy. Large language models train on data that may include personal data, sometimes incidentally and sometimes to achieve important purposes. Engineering techniques such as "model unlearning" are still in the early stages of development (whereas others, such as output filters, are more mature). Across all stages of training, employing techniques to appropriately record the provenance of data is vital for transparency and explainability, as well as assuring data quality.

Data minimization is contextual, and some contexts require personal, and even sensitive, data. For example, training models that guide the operation of self-driving cars require precise location data (down to 1 cm), as the stakes are high due to possible human harm. Rather than set absolute thresholds for the amount of data that should be collected, it is important to consider the specificities of the intended use cases and potential life cycle approaches.

Should transparency and explainability be mandated? Of key importance is who will be defining what we will be explaining, what does the consumer want to know, and how can the information provided be made digestible and relevant to them? The answers may differ depending on the audience (e.g., consumers, regulators, and business deployers). Consider instituting privacy "nutrition" labels with important information or publishing model or system cards alongside any product releases. (See, for example, Internet Safety Labs.)

Think about actual harm instead of hypothetical harm. Effective risk assessment and mitigation for generative AI require being rigorous in building typologies of potential harms and distinguishing between theoretical or low-likelihood risks and those that are actually being experienced. It is crucial to work with the engineers to build effective safeguards into design from the outset.

Privacy expectations vary across different individuals and contexts. What might be a necessary feature for one person or in one context may be problematic or a risk factor for another user. Special care should be taken to address risks associated with youth and children.

Standards for privacy in AI have a valuable role to play. Privacy engineers would benefit from standards to work toward with respect to engineering privacy into AI models, as well as to establishing provenance.

Privacy engineers in organizations deploying AI are facing the challenge of assuring the quality of models engineered elsewhere. Because deployers are often using models developed by engineers in another organization (e.g., third-party tools or applications, open-source software), they don't necessarily know what was done before with regard to data management at this foundational level. Transparency from the developer is vital to ensure that the development practices are consistent with your company's privacy policy.

Discussion on "A Balancing Act: How PETs can deliver a win-win for personal privacy and corporate goals"

Align privacy-enhancing technologies (PETs) with business objectives. Start by aligning internally on the problem you want to solve, identifying values and pain points, and investing in privacy for the long term.

The cost of PET technologies is a key barrier, although many PETs are becoming more affordable to deploy at scale. Consider incentives for implementation other than government regulations. It is essential to rigorously define the business case for use of PETs and to select the most appropriate PET to deploy to address the specific needs of the scenario in question. Include the cost of addressing privacy breaches when evaluating the cost of PETs.

Retrofitting privacy is difficult. In many organizations, there are often old systems and new systems working together, many created before the development of current PETs. However, experience has found that it is optimal to implement privacy engineering with PETs from the ground up whenever possible.

Implementing PETs may initially require engineering tradeoffs. Engineers may need to determine what are the available tools and what levels of risk are acceptable, as it is often difficult to generate high-performance models with some PETs. Additionally, various types of systems handle customer data, and different PETs are required for each system. However, privacy by design highlights a positive approach to achieve full functionality (PdB Principle #4).

Look for design inspiration from other sectors. Consider borrowing tools from the field of data security and adapt them for privacy.

Design a more objective test or approach to PETs deployment. Partner with different data teams and test some PETs, noting the challenges, benefits, and disadvantages of each in order to design a metric for proof of concept. Also investigate secondary reviews: What do users like in privacy settings and what changes would they like added? If you can't get user validation, consider other types of validation or metrics.

Sometimes business confidentiality considerations limit our ability to talk about our PETs deployments externally. Because of this, we lack forums to discuss challenges from a technical perspective. In what ways can best practices be shared?

Could the use of synthetic data be mandated? Synthetic data can help to minimize risk in use cases where personal data would otherwise be needed. However, good quality synthetic data comes at a cost, which may be prohibitive for smaller companies.

PETs may be used to address the challenge of keeping data safe when control passes from one organization to another. PETs can be especially useful for data security when data is transferred out of a secure environment into another.

Best practices for managing regulatory guidelines and approval. If you have a PET that works, talk to a regulator in advance and develop a relationship. If you wait to convey vital information after implementation at their request, you may miss the opportunity to avoid potential pitfalls.

Enhancing compliance and adoption. Consider how to pitch PETs internally as problem solvers. Develop technical metrics. Work with external validators, such as advocacy groups, to help you raise awareness of the purposes and benefits of PETs and make the case for the value they can bring.

Outside organizations such as IEEE could assist by defining terms and guidelines for consumers. Users might not be familiar with PETs, but they care about being safe. The use of technical lingo makes it harder for users to participate in conversations on privacy. Consider how privacy concepts could be better explained to users effectively in real time, e.g., pop-up explanatory notices.

Education about privacy should happen at multiple levels. Privacy needs to be part of the engineering curriculum and discussions should also begin in earlier grades. Canada has embraced an early education model. The IEEE Digital Privacy Initiative is developing curriculum guidelines to assist universities in training the next generation of privacy professionals. In addition, privacy courses designed in tandem with key universities are available for continuing education (CE) credits.

Additional Resources

Publications

CIPL, *Reconciling AI with the Data Minimization Principle: Bridging the Innovation and Privacy Gap* (forthcoming).

CIPL, *Rethinking Sensitive Data in the Age of AI*, September 2025.

CIPL, *Privacy-Enhancing and Privacy-Preserving Technologies in AI: Enabling Data Use and Operationalizing Privacy by Design and Default*, March 2025.

CIPL, *Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework*, February 2024.

CIPL, *Privacy-Enhancing and Privacy-Preserving Technologies: Understanding the Role of PETs and PPTs in the Digital Age*, December 2023.

IEEE Digital Privacy Model, 2023.

IEEE, IEEE Code of Ethics, June 2020.

CIPL, *What Good and Effective Data Privacy Accountability Looks Like: Mapping Organisations' Practices to the CIPL Accountability Framework*, May 2020.

IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, *Ethically Aligned Design*, First edition, 2019.

Information and Privacy Commissioner of Ontario, *Privacy by Design*, January 2018.

Websites

Centre for Information Policy Leadership (CIPL)

IEEE Digital Privacy Initiative

IEEE Privacy Course Program

IEEE Standards

IEEE Education

Internet Safety Labs

Thank you to our participating organizations:

Abaxx Technologies

Adobe

Airbnb

AppCensus

Apple

Axon Enterprise

Brave Software

CableLabs

Cisco Systems, Inc.

Consumer Reports

DoorDash

Electronic Frontier Foundation

Empower Privacy

Hofstra University

Internet Safety Labs

JLINC Labs

Mastercard

Meta

Microsoft

PayPal

International Computer Science Institute (ICSI), UC Berkeley

Rivian VW Tech

Salesforce

SAP

The Walt Disney Company

TrustArc

University of Illinois at Urbana Champaign

Whirly Labs