From Principles to Operational Reality: Measuring and Governing Responsible AI at Scale

Takeaways from Roundtable held in conjunction with the IEEE Symposium on Privacy Expectations (ISoPE)

1 October 2026

Takeaways from Roundtable

From Principles to Operational Reality: Measuring and Governing Responsible AI at Scale

HELD 01 OCTOBER 2026
NEW YORK, NEW YORK

As organizational use of AI matures, the focus is increasingly shifting from high-level responsible AI principles to operational execution and measurable outcomes. Regulators now expect organizations not only to articulate commitments to accuracy, non-discrimination, transparency, and accountability, but to demonstrate how these principles are measured, monitored, and enforced in practice.

On October 1, 2026, the Centre for Information Policy Leadership (CIPL) joined the IEEE Digital Privacy Initiative to convene a roundtable in conjunction with the IEEE Symposium on Privacy Expectations (ISoPE) in New York City. The roundtable discussed how organizations are translating AI principles into concrete governance structures, technical controls, and monitoring processes across the AI lifecycle.

Participants included senior professionals from CIPL member companies and the IEEE. The event was conducted under the Chatham House Rule to foster open and candid dialogue.

CIPL is pleased to share the following takeaways.

- **Organizations are still grappling with how to operationalize responsible AI principles.** The challenge is translating broad commitments into clearly assigned responsibilities, enforceable system constraints, and continuous oversight.
- **There's too much reliance on traditional measures of accuracy (such as measuring false positives) which principally rely on classification of data sets.** These measures do not fully capture whether an AI system's behavior is appropriate for its intended use, particularly when evaluating open-ended outputs or autonomous actions.
- **We need to be open to experimenting with new methods of measuring accuracy.** Such experimentation could include evaluations that assess performance against task-specific objectives, account for uncertainty, and test behavior under changing real-world conditions.
- **Most answers are contextual, but that's not even part of the conversation yet.** Evaluations need to account for the intended use, affected stakeholders, and consequences of error rather than treating correctness as a context-independent property.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

- **Risk tolerance comes down to "vibes," i.e., our gut says it's okay to move forward.** If we want to move forward in a meaningful way, we need provable, mathematical processes, which run against fundamental business processes. Moving beyond intuition would require explicit acceptance criteria and evidence-based approval decisions, potentially creating tension with business expectations for speed and flexibility.
- **Some are exploring "dependent type theory" which seeks to reduce subjectivity.** Dependent type theory is a foundational framework in mathematical logic and theoretical computer science where types can depend on values. By allowing types to encode logical propositions, this approach could help make certain requirements explicit and formally verifiable, although selecting those requirements still involves judgment. Dependent type theory is being evaluated as an antidote to "vibe-based" risk management and subjective oversight.
- **Some would like to see regulatory guidelines, but others vociferously objected, as guidelines can be incredibly prescriptive.** The tension is between providing enough clarity to guide implementation and retaining the flexibility to accommodate different use cases and evolving technologies.
- **In the agentic context, credentialing agents is an insufficient solution because credentials can be compromised.** Establishing an agent's identity needs to be complemented by controls over its permissions, ongoing scrutiny of its actions, and mechanisms to revoke access.
- **There is no mathematical certainty in agentic technologies today.** This concern highlights the distinction between proving specific properties under defined assumptions and guaranteeing the behavior of an entire agentic system in an open-ended environment.
- **Internally, companies are talking proactively, but acting reactively in light of constantly changing regulatory goal posts.** The resulting challenge is to establish durable governance practices while preserving the ability to adjust controls as regulatory expectations evolve.
- **Engineers will need to partner with academia, social scientists, and others in the humanities to craft appropriate solutions.** These collaborations can help connect technical safeguards with the ethical judgments, institutional contexts, and human consequences that engineering methods alone cannot resolve.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.

Who We Are

The Centre for Information Policy Leadership (CIPL) is a global privacy and data policy think tank within the Hunton law firm that is financially supported by the firm, 70+ member companies that are leaders in key sectors of the global economy, and other private and public sector stakeholders through consulting and advisory projects. CIPL's mission is to engage in thought leadership and develop best practices to ensure the responsible and beneficial use of data in the modern information age. CIPL's work facilitates constructive engagement between business leaders, data governance and security professionals, regulators, and policymakers around the world. For more information, please see CIPL's website at https://www.informationpolicycentre.com/. Nothing in this document should be construed as representing the views of any individual CIPL member company or of the Hunton law firm. This document is not designed to be and should not be taken as legal advice.

The IEEE Digital Privacy Initiative was approved under IEEE Future Directions in January 2022 to cultivate emerging technologies to address digital privacy for individuals. The Initiative focuses on a user-centric perspective – looking at the digital privacy needs of the individuals rather than the security of data, products, and organizations – such as providing individuals with user-enabled privacy controls and promoting privacy at the outset of product and service lifecycles.

Security and privacy are interdependent, and privacy cannot be achieved without security. However, these terms are not synonymous. Security focuses on protective measures to prevent malicious threats and protect data from unauthorized access. Security also protects the interests of both individuals and enterprises. Privacy focuses on ownership establishment, control one has over his or her own personal information, and what/where/how the personal information is collected, processed, shared, and used.

The IEEE Digital Privacy Initiative envisions a future in which the capability exists to enable any individual around the world to privately maintain presence, data, identity, and dignity online. To help achieve this vision, the Initiative seeks the following goals:

- bring the voice of technologists to the digital privacy conversation, incorporating a holistic approach to address privacy that also includes economic, legal, and social perspectives
- facilitate cross-disciplinary collaboration to advance research, promote standardization and best practices, and create tools and capabilities to support the privacy needs of individuals, and
- coordinate efforts across and beyond IEEE with a multicultural lens that are working on different dimensions of digital privacy.

Copyright © 2026 by the Centre for Information Policy Leadership at Hunton Andrews Kurth LLP.