

Centre for Information Policy Leadership
HUNTON

AI for Compliance

Advancing Accountability and Trust in Data and Digital Governance

November 2026

Table of Contents

Introduction: Why AI for Compliance, Accountability, and Trust? 3

Project Overview 5

Key Insights and Best Practices 6

Foundational Concepts and Considerations for AI for Compliance 8

AI for Compliance: Opportunities & Considerations 11

1. Scaling and Transforming Compliance Operations and Outcomes 11
2. Improving Accountability and Transparency 13
3. Enhancing Risk-Identification and Decision Making 14
4. Standardizing Policies, Controls, and Safeguards 16
5. Enabling Real-time Audits 17
6. Augmenting Human Expertise 18

Considerations for Regulators and Policymakers 20

Conclusion 22

Appendix I: Use Cases 23

Appendix II: Key Considerations for Designing and Deploying Automated Compliance Tools 26

Introduction: Why AI for Compliance, Accountability, and Trust?

Compliance and product teams face similar challenges in organizations of all sizes—rapidly evolving and complex digital regulation, business pressure to drive faster product development and decision making, increased external expectations from regulators to demonstrate accountability, and limited resources and time. This is particularly true for digital governance and compliance matters in related digital disciplines, such as privacy, data, AI, safety, and cybersecurity.

Externally, privacy, data, and AI governance teams play a key role in building digital trust and confidence. As data takes on a more central role in the development and delivery of products and services, expectations for accountability and transparency from customers, business partners, investors, regulators, and the public will continue to steadily grow. Also, internally, these teams are central to the development of safeguards and accountability measures, as well as enablers of responsible data-driven innovation. Finally, with increasing AI capabilities and growing enterprise adoption of AI, compliance teams must ensure that internal AI uses comply with organizational principles, policies and procedures.

Within this context, compliance teams are increasingly exploring the use of AI-powered tools to improve compliance outcomes for individuals and the organization as well as effectively monitor and oversee the AI-powered systems, models, and insights that are driving business decisions today. Privacy and data management compliance tasks (e.g., records keeping, data lineage and mapping, purpose limitation, retention controls, risk assessments, third-party diligence and oversight, data security measures, individual rights and complaint handling processes) naturally fit into an automated stack and can benefit from AI-powered technologies due to their repetitive nature and the large volumes of data that must be analyzed.

Importantly, when deployed responsibly, AI for compliance solutions present an opportunity for organizations to:

- Improve compliance outcomes and accountability by extending compliance coverage and removing bottlenecks across all phases of the compliance lifecycle, particularly as compliance has historically relied on manual processes that are time- and resource-intensive and difficult to scale;
- Drive faster, repeatable, standardized, and more accurate decision-making; and
- Free up compliance professionals' resources that may be better spent on the most strategic, high-value work and tasks.

While organizations stand to gain many benefits from the automation of and adoption of AI for compliance, there will be ongoing risks and challenges associated with early adoption of any technology and the internal change management required. For example:

- Automating compliance can raise concerns about human oversight of AI-powered tools, explainability, trust, accuracy, and even the future of the compliance profession.
- There are regulatory, technical, and reputational questions about what level of automation is appropriate for any given task or process.
- Compliance teams are generally composed of legal and compliance professionals who can be more risk-averse to the adoption of AI-powered tools and may require additional upskilling in technical understanding to deploy these tools to their greatest potential.
- Regulatory expectations and attitudes towards AI tools for compliance and accountability may vary across different countries and jurisdictions.

The adoption of AI is at a critical crossroad for compliance teams, not only presenting profound opportunities if done responsibly but also raising legitimate challenges that must be effectively managed. This represents a shift in many organizations' focus and priorities—from focusing solely on AI governance to a much broader focus on governance by AI, or *AI for Governance*.

For all these reasons, the Centre for Information Policy Leadership (CIPL) has launched a new global project, AI for Compliance, that sets out to better understand the benefits, challenges, and industry best practices for deploying AI for compliance, accountability, and trust. This report details the first project findings, emerging industry best practices, industry use cases, and recommendations for stakeholders from industry, policymakers and regulatory bodies. These considerations are intended to promote the trustworthy and responsible adoption of AI for compliance through improved accountability and empowered humans in the loop. Considerations are also intended to inform future research, industry benchmarking, and multistakeholder dialogue on the appropriate policy, legal and technical considerations for automated compliance.

This project builds on CIPL's ongoing work on organizational accountability in data protection and broader digital and data policy, ongoing for more than 25 years. In 2018, CIPL identified how organizational accountability can serve as a key building block not only for effective data protection, but also for the responsible governance of AI. CIPL's projects and thought leadership seeks to advance accountability in AI governance within operational reality. AI for compliance is a natural next step of CIPL's accountability work and an essential method for companies to effectively deliver all elements of CIPL's Accountability Framework by implementing the latest AI-powered technology and tools.

The AI for Compliance project also draws insights from CIPL's 2024 report *Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework*, where CIPL collected and identified organizational best practices from leading companies' AI governance programs, as well as CIPL's ongoing efforts to promote accountable governance in the era of data-driven innovation.

Project Overview

This first report is based on discussions that took place during two live workshops that CIPL co-hosted in partnership with Meta's Trust, Transparency and Control Labs (TTC Labs), using their Design Jam methodology, as well as CIPL's independent research, webinar and interviews with CIPL member companies.

During the AI for Compliance project, CIPL and TTC Labs co-facilitated Design Jam workshops with stakeholders from industry and academia. These interactive workshops were led by multidisciplinary experts and explored key themes underpinning the development and adoption of AI-powered tools for compliance, including:

- Foundational AI concepts and capabilities;
- The opportunities and risks concerning potential industry-wide use cases of AI for compliance;
- The role of humans in processes and systems embedded with AI; and
- How can AI support organizational accountability.

Project participants represented the following sectors:

- Academia;
- Automotive;
- Banking and finance;
- Consulting;
- Entertainment;
- Legal;
- Non-profit;
- Pharmaceuticals;
- Social media
- Technology; and
- Retail.

Project participants represented a wide range of experiences, perspectives, and maturity levels concerning the adoption and role of AI in compliance tasks and processes. Their roles included policy, compliance, legal, product, technical, and academic, and they represent global organizations that reach billions of individuals. The workshops were held under the Chatham House Rule to encourage open discussion and facilitate trust between project participants.

Key Insights and Best Practices

1. There is **high enthusiasm among compliance professionals for automation**, but wide-scale adoption hinges on literacy, trust, explainability, and governance. Rapidly evolving regulatory landscapes, increasing digital regulation and compliance requirements, and business pressure to speed up product development while maintaining high levels of legal and regulatory compliance are all motivating compliance professionals to turn to automated tools.
2. The appetite for automated compliance is highly dependent on the **use case, risk profile, and organizational context** of the tool. There is no one-size-fits-all approach. Organizations are progressing at different speeds, with adoption shaped by **regulatory exposure, internal maturity, and the nature of the compliance task.** An AI compliance application can be low-risk in one organization but high-risk in another, depending on various factors such as the industry, related regulations and the sensitivity of the underlying data processing and use case. Some regulated industries, such as financial services and pharmaceuticals, as well as technology companies, are naturally more open and ready to invest in AI for Compliance tools. Indeed, an article published in the MIT Computational Law Report notes that regulatory complexity in the financial, insurance and legal sectors demands AI based solutions.[1] Similarly, companies that have already developed and implemented mature compliance programs, either as a result of GDPR, other national laws or standards, or regulatory consent decrees (such as those issued by the US FTC) across various digital disciplines, including data privacy, cybersecurity, AI, competition, and safety, may be better positioned to develop and deploy AI for Compliance tools. This is because they oftentimes already possess the necessary "know-how", data, precedents, and evidence of their established compliance programs and controls. These elements establish a compliance corpus of best practices that can be used to train and develop AI tools to automate and scale compliance activities and tasks.
3. **Compliance professionals are in a unique position** whereby they are exploring the value of AI technologies within their remit as well as managing the risk that these same technologies present to the business. Unlike other business functions, e.g. marketing or sales, the compliance function's relationship with AI-enabled automation is fundamentally rooted in managing and mitigating risk while ensuring legal compliance.
4. **Compliance professionals must remain accountable for their AI-automated processes and outcomes.** Compliance professionals and organizations must remain fully accountable for outcomes, regardless of whether automation is involved in the process. This requires organizations to **intentionally design with accountability in mind from the outset**, embedding mechanisms that empower and incentivize compliance professionals to **evaluate, verify, and challenge AI outputs**, rather than passively accepting them.

1 Sobkowski, M., & Karapetyan, G. (2025). The Dawn of a New Era of Compliance: Automated Compliance Verification and Enforcement. MIT Computational Law Report.

5. **Data quality is the foundation of trustworthy automation**. The effectiveness of AI-powered compliance tools depends fundamentally on the quality, relevance, and integrity of underlying data and sources. Without strong data foundations, even advanced systems will produce unreliable outputs.
6. **AI literacy is essential to responsible adoption.** Compliance professionals must understand how AI systems work, their limitations, and how to interrogate their outputs, ensuring that automation enhances, rather than undermines, professional judgment and accountable business practices.
7. **Adoption should be iterative and risk-based.** Organizations should begin with low-risk, high-volume use cases, using these as a foundation to build trust, refine controls, and develop internal capabilities. Over time, automation can scale alongside increased organizational maturity, stronger governance, and improved AI literacy.
8. **AI should augment, not replace, human judgment.** AI is most effective as a compliance partner, not a substitute. It enhances scale, consistency, and efficiency, but human expertise remains essential for interpretation, ethical judgment, and high-stakes decisions. The future of compliance is defined by a human–AI partnership model, not full automation.
9. **Automation expands coverage and strengthens accountability.** Compliance teams are turning to AI not just to save time, but to also extend their coverage across more processes, risks, and business activities. By enabling broader and more consistent oversight, automation can strengthen organizational accountability and improve compliance outcomes at scale.

Foundational Concepts and Considerations for AI for Compliance

While organizational pressures to move faster and at scale are motivating compliance teams to adopt automation in their work, improved compliance outcomes are demonstrating the strategic value of investing in AI-tools and the necessary governance involved with their deployment. At the same time, the nature of compliance work, ensuring that an organization operates legally, ethically, and in alignment with both external and internal governance principles, means that compliance professionals need to be able to trust AI-powered tools in order to rely on them and realize their full potential.

AI Literacy

AI literacy is a foundational prerequisite and critical enabler for deploying AI-powered compliance tools. It is the connective tissue between AI capability and accountable deployment. Organizations must invest in AI literacy to enable compliance teams to design, experiment, effectively develop, and deploy automated compliance tools. Organizations that fail to prioritize AI literacy can get stuck in a "AI is useful but hard to deploy" stage and risk disadvantaging their compliance programs. Additionally, failure to upskill compliance professionals can result in disempowered human oversight of AI systems and models and the risk of AI misuse.

Stages of Automation

STAGE 0 | Human Execution
Tasks are completed by humans without AI support.

STAGE 1 | AI Assistance: AI offers suggestions.
Humans decide what is useful and relevant.

STAGE 2 | AI Proposes
AI proposes recommendations. Humans accept or reject.

STAGE 3 | AI with Guardrails
AI completes routine tasks with defined boundaries.

STAGE 4 | AI Runs Most
AI handles most of a task end-to-end. Humans handle only novel, complex or contentious situations that meet material risk thresholds.

STAGE 5 | Full Automation
AI runs all processes related to a task. Humans design, train, audit and evolve the system.

AI-powered tools vary in terms of their level of automation for any given task or process. An automated task does not necessarily mean Stage 5 full automation. Many organizations maintain and reference a 5- or 4-stage autonomy framework, such as the one above, to explain AI capabilities and guide internal considerations and controls. Importantly, an autonomy framework is most useful when each stage includes actionable considerations related to risk and controls. For example, the autonomy framework can help organizations establish proper human-in-the-loop controls or triage whether a use case needs further review depending on the context and impact of the underlying task. Further, certain laws and regulations will prevent Stage 5 full automation or even stage 4 automation, depending on interpretations and risk thresholds. Many project participants agreed that individual compliance responsibilities and tasks that do not directly impact individual rights are easier to fully automate, while compliance frameworks and tasks that involve material risks and individual rights will be more difficult to fully automate.

An autonomy framework, with characteristics that delineate between each stage, can help guide a compliance team design AI-powered tools with the appropriate level of automation depending on the task. Project participants highlighted the need to gradually build proper AI for compliance governance and controls to increase from one stage of automation to the next.

Recommendations from the Field

Establishing AI governance policies, providing access to enterprise AI tools, and offering AI literacy trainings is one thing. Organizations need to also take steps to enable their workforce, particularly compliance teams, to use the tools responsibly and effectively. One organization launched a workforce AI training program where certain individuals are trained and enabled with the resources and skills to develop AI agents. These individuals are empowered to then support their colleagues with agentic AI solutions. Another organization embeds technical and security staff within their legal and compliance teams to help explain the technical concepts that could be relevant to their work as well as the security risks that should be considered along the way. Exposing legal and compliance teams to the relevant technical concepts, skills, and considerations is the foundation for effective AI for compliance programs.

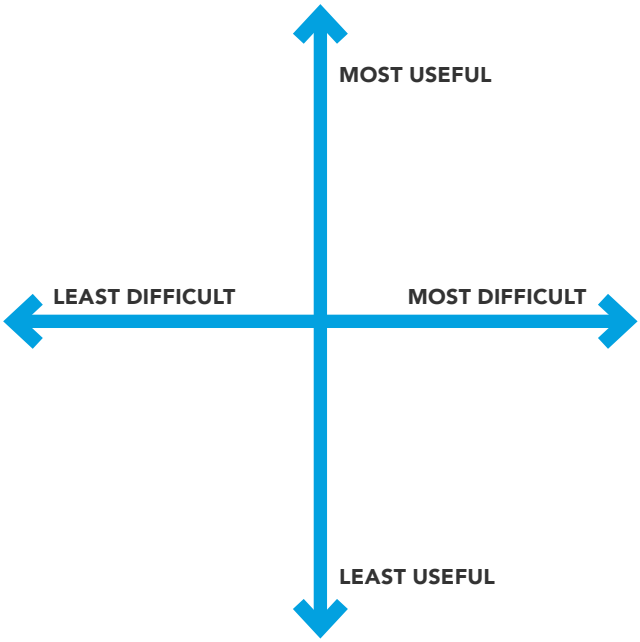
Data Quality

AI-powered compliance tools will only be as effective as the underlying data sets and sources that they are trained on. Organizations will need high-quality and trustworthy data sets that thoroughly reflect historic compliance decisions and processes to develop AI systems and models that can produce reliable outputs. High-quality data sets must have accurate and up-to-date data that effectively reflects the regulatory landscape and organizational policies.

It is also critical to continuously assess the decision-making of AI systems to mitigate or correct for potentially biased behavior or outdated signals. This can be made more difficult when AI models and systems face edge cases and novel situations, as the measurement of bias or risk may not be clear. In these circumstances, compliance professionals must have a deep understanding of the technology as well as be empowered to make a decision and effectively intervene.

AI Capabilities x Value

A useful exercise for compliance teams considering automation entails a survey of all relevant tasks and potential AI capabilities. Compliance tasks, or elements of tasks, that align with AI capabilities should then be mapped on a matrix depending on usefulness and difficulty. Visualizing and understanding the level of usefulness or impact and the level of ease related to deploying such an AI tool will help organizations create achievable automation goals and balance them against competing resource and capacity limitations or demands. For example, a less useful task that is difficult to deploy may not be worth prioritization while a very useful task that is easy to deploy could be a helpful test that helps upskill the compliance team and earn trust in the AI-tool. Organizations can also track elements that make tasks and AI capabilities useful or difficult to deploy to encourage documentation and assessment.



Data Protection and Privacy Considerations

In consideration of the AI capabilities x value matrix, the compliance function within an organization should also consider the concrete data protection and privacy tasks and responsibilities when assessing opportunities for automation. Legal requirements such as access, correction, and deletion requests, complaints, legitimate interests balancing, purpose identification and tracking, records of processing, and assessments may benefit from varying degrees of automation, especially where an AI tool that performs reliably well can improve outcomes with faster response times, improved accuracy, and a higher level of standardization. The value of time saved on such tasks can support compliance teams in focusing their limited time and resource of complex issues that meet material risk thresholds. Accountability-based tasks and controls, such as reporting, verification and monitoring, training and awareness, breach notification, and risk assessment often rely on complex frameworks and have far-reaching impact and consequences on organizational decision-making. While these tasks will also benefit greatly from automation, the benefits may be more complicated and difficult to realize.

AI for Compliance: Opportunities & Considerations

AI has the power to transform the compliance role from a reactive to proactive function. Historically, compliance teams respond to regulatory developments during product reviews, establish documentation post-hoc, after decisions have been made, and scramble to keep up with business decisions and needs. Reactive compliance teams can be perceived as bottlenecks that slow down business operations. AI-powered tools, deployed within a broader culture of organizational accountability, can transform a compliance team into a proactive function that is integrated into daily operations. The proactive compliance team can spend less time on monotonous, repetitive tasks and free up time for the most high-value strategic work.

Throughout the AI for compliance project, participants considered how AI-powered tools can make them more effective in their role—more accountable and more thorough, with broader impact and reach across their organizations. One participant considered how automation and AI-tools can support better human decision-making. Many legal and compliance teams operate under growing pressure with an ever-expanding regulatory landscape, increasing transaction volumes, shrinking timelines, and limited headcount. This operational reality is motivating compliance professionals to explore how AI-powered tools can enable them to meet growing demands with speed, quality and at scale. Project participants discussed "AI-enablement" and "AI-augmentation", they want tools and processes that can equip them to make better decisions at higher volumes rather than replace their decision-making. While the maturity level of AI for Compliance programs varies widely across participating organizations, the following opportunities resonated with diverse industry stakeholders.

1. Scaling and Transforming Compliance Operations and Outcomes

Participants shared tangible benefits, including efficiency and productivity gains, improved consistency and standardization, and overall strategic business value gains. AI-powered tools can reduce voluminous, time-intensive tasks, in some cases from days to minutes.

For example, one organization built an internal automated cookie compliance tool that completely transformed the amount of time the legal team had to spend reviewing dense spreadsheets. The automated cookie tool suggests categories and controls, maintains a record of all reviews, and the legal team reviews the cookie categorization recommendations before anything is finalized. This tool freed up valuable capacity for more complex tasks and evolved a demanding and human error-prone task.

Another organization has experienced greater than a 60% increase in data subject access requests (DSARs) in some jurisdictions, in part due to the use of generative AI to create and manage requests on behalf of the data subject, which has placed added volume and pressure on the compliance team. Automation tools have greatly helped this

organization manage increasing DSAR volumes in a timely manner, thereby reducing legal risk, while also improving the accuracy and thoroughness of responses. Complex DSARs are routed for human processing and review.

Is Full Automation Possible?

The stage of automation appropriate for a given compliance task is highly dependent on the use case. One organization maintains human review of even low-risk, repetitive tasks that have been automated until it achieves sufficient confidence in automation. The organization's goal is for an automated task to be completed with better accuracy than by humans. Once this threshold is met, human review is relaxed but never completely eliminated as humans maintain audit and design tasks. AI systems are continuously monitored, outputs are verified, and models are fine-tuned before further automation is designed into the system.

Routine, well-defined tasks, like baseline DSAR management, records of processing and data flow mapping, or third party due diligence processes are well suited for high stages of automation. As are tedious, repetitive tasks that benefit from standardized approaches, like documentation and contract review. However, AI-powered systems must be designed to recognize novel, high-risk, or ethically complex cases, which should be routed to compliance professionals.

Considerations

- An overwhelming majority of project participants agreed that AI-adoption without proper safeguards and controls is a bigger risk than no AI at all. Safeguards and controls need to account not just for systematic errors and security risks, but also for novel, edge, and ethically-complex patterns and circumstances.
- Overreliance on high-levels of automation can create new kinds of risks, including compliance professionals ill-equipped to audit AI systems and critically assess subtle, evolving threats, if necessary. Organizations should carefully scale automation in compliance functions while managing literacy, governance and impact.
- AI-powered compliance tools are most effective when designed in a collaborative manner across an organization, for example between compliance professionals and their technical and security peers. Many project participants described a decentralized, bottom-up approach to AI-adoption while others described receiving unsolicited tools from leadership or technical peers that did not reflect team priorities. Developing automated compliance tools in a completely unstructured and decentralized environment can delay the transformative potential that AI tools can provide or at worse expose the organization to new levels of risk, for example from prompt injection attacks. A decentralized approach can also lead to the development of overlapping tools that are misaligned and resource-draining. Meanwhile, an organization that develops and adopts AI tools without input from compliance professionals risks overlooking priorities and misplacing resources. A coordinated, centralized approach, where compliance professionals are involved in the design process, helps ensure that AI-powered tools properly balance priorities, risks, and the potential for the highest value add to the organization.

Case Study

One organization has embedded a dedicated data scientist/technologist within its privacy, legal, and compliance team to support the automation of data protection requirements and program controls. This cross-functional approach combines legal and compliance expertise with technical capabilities to identify opportunities for automation and translate regulatory requirements into operational tools and workflows.

2. Improving Accountability and Transparency

Organizational accountability in data protection operates as the bridge between legal requirements and real-world, effective protection of individuals' data. It is a framework that allows organizations to demonstrate that they do what they say will do, as well as what the law and regulations require of them. CIPL's Accountability Framework shows how accountability is operationalized through seven elements: 1) leadership and oversight; 2) risk assessment; 3) policies and procedures; 4) transparency; 5) training and awareness; 6) monitoring and verification; and 7) response and enforcement. These elements help organizations turn abstract legal principles (like fairness, purpose limitation, etc.) into policies and procedures, risk assessments, controls, and safeguards.

This work is time intensive and hinges on many individuals and teams across an organization, including legal, compliance and audit, product, and business teams (often organized as first, second and third lines of defense), who must translate policies and procedures into operational reality. AI-powered compliance tools can help standardize the application of the accountability elements and integrate cross-function controls and tasks. Automation also reduces the likelihood of human error, such as incorrect risk scoring, insufficient documentation, or inadequate safeguards. Increased capacity within a compliance team can also expand the team's reach within an organization and elevate the baseline level of accountability reinforcing organizational decision-making.

Case Study

One organization has integrated the privacy, risk, compliance, information security and audit functions into a single "trust" function. This evolution has allowed the organization to simplify governance, accelerate decision-making, strengthen accountability and comprehensively assess risk. The new trust function used AI to completely redesign and automate the organization's security and compliance process, improving review time and risk evaluation. As part of the security and compliance process, any new technology use case must be evaluated for privacy-by-design, security-by-design and other compliance related considerations. Employees are trained to input detailed use case approval requests, including relevant characteristics and potential risks, and provided with immediate guidance. Requests are triaged with low-risk use cases granted immediate approval and other requests assessed against a medium-, high- or material-risk framework, with each level requiring more in-depth human review. The entire process is automated and fully traceable. This has allowed the organization to improve accountability, internal and external transparency and reduce human-driven assessment volumes.

Several project participants emphasized that AI-powered compliance tools have strengthened their accountability programs through comprehensive audit trails, continuous monitoring and reporting, and enhanced visibility into decision-making processes. For example, across many participating organizations, automation tools that map and track internal data inventories and data flows have evolved the repetitive and labor-intensive nature of completing records of processing activities (ROPAs) resulting in less errors and more complete records that reflect up-to-date information. These benefits also improve transparency, making accountability demonstrable to individuals, regulators, and other stakeholders.

Considerations

- Organizations must define who is responsible for compliance decisions when AI-powered models and systems are involved. This will reduce the risk of over-reliance on AI models and systems and ensure that compliance professionals and organizations remain accountable for decisions.

- While generative AI tools can produce audit trails, certain AI models, such as neural-based AI models, produce outputs that are not easily explainable. Organizations must consider the use case of the AI-powered tool and its impact on material risk thresholds and individual rights to determine explainability requirements and limitations.
- Just because something can be automated does not mean it should be automated. Various considerations, including reputational and regulatory consequences, should determine whether a compliance-related task should be automated or not. Some organizations intentionally limit AI-powered tooling in certain compliance tasks and processes and retain 100% human-decision-making. For example, one organization determined that certain matters related to employee data are not worth automating due to the sensitivity of the data and the underlying litigation risk. Another organization determined that there are too many cultural and reputational risks, as well as questionable efficiency gains, associated with fully automated hiring and maintained human decision-making in key recruitment and hiring processes while incorporating automation to augment human decision-making.
- Compliance professionals are concerned about the tradeoffs between transparency and liability. Some project participants raised concerns that AI-powered systems that generate comprehensive audit trails may produce regulatory and litigation exposure risks for documentation that should otherwise be protected by legal privileges. Organizations will need to work carefully with their legal counsel to develop policies and procedures that ensure legal privileges are preserved when and where appropriate.

3. Enhancing Risk-Identification and Decision Making

The risk management cycle is an important element of operationalizing data protection and data governance related requirements. In practice, risk management is a continuous and time-intensive compliance process that is constantly evolving based on external and internal developments. AI-powered systems can support sharper, faster risk analysis and earlier identification of issues when the systems are effectively integrated across an organization's operations. They can also allow for repeatable, more consistent risk assessments and scoring processes that can help minimize the risk of human bias and subjectivity.

Many project participants are exploring automation to reduce the time and effort it takes to complete risk management related tasks, such as data protection impact assessments (DPIAs) and privacy impact assessments (PIAs). In addition to completing these time-intensive tasks faster, compliance professionals want to standardize how risk is assessed and scored based on likelihood and severity, when appropriate, and extend a baseline level of compliance coverage to all organizational activities that require risk assessments. Finally, enhanced risk assessment outcomes can inform deeper understanding of the overall risk position of the organization and enable more accurate, real-time reporting to upper-level management, audit and risk committees, and the board and C-suite..

Case Study

One organization's risk review process assesses a broad range of risks (data use, security, youth controls, etc.) across new and updated products, features, and data practices, conducting an average of 1,800 reviews every month. It enables the organization to understand any risks associated with a launch and assess and apply the product requirements needed to mitigate those risks. Over years of launching products into global markets with unique regulations and expectations, the organization has developed a corpus of commonly applied requirements.

Traditionally, assessing each launch for the necessary requirements depended on the product owner answering dozens of questions per risk category (e.g. sensitive data use, third-party integrations, etc).

The right requirements from the corpus were then applied using predefined rules. While this approach effectively guides product owners through a rigorous risk assessment, it is difficult to scale as the complexity of products and the regulatory landscape increase over time. A product owner could face hundreds of questions for a single feature launch, including for rare risks that are not relevant to their launch. Whilst each product area is responsible for effectively managing their product's risk, it's not reasonable to expect every product owner at the company to become a global risk expert.

The challenge: applying the right requirements, at scale. The organization now automates requirements application for each review. Using AI and automation, Risk Review gathers context directly from the artifacts product teams already produce, such as code changes, product requirements documents, and configurations. Risk Review can then automatically identify and apply the right requirements to each project (and trigger for specialized human input for novel or complex decisions). Because every requirement is evaluated against every launch, this automated workflow not only reduces the burden of expertise on the product owner, but also enhances assessment rigor and consistency.

However, applying the right requirements is only part of the task; each one then has to be evidenced to prove it was implemented correctly in the product. In this new AI-driven workflow, evidence for applicable requirements is gathered and pre-filled from product team sources. An independent verification step then reviews the evidence, makes an explicit determination (satisfied or insufficient) and outputs structured reasoning explaining why. The accountable product owner reviews the evidence, reasoning, and determinations, can iterate or edit where something is insufficient, and must explicitly confirm before submission—after which the review remains subject to independent human audit.

Importantly, as the organization applies AI and automation to specific steps in the Risk Review workflow, every launch is still assessed for the same risks and requirements; product teams remain accountable for their decisions as the first line of defense; and audit, oversight, and remediation processes are unchanged. Humans remain the architects and overseers: designing the rules, setting the standards, validating performance, and retaining override authority.

AI-powered risk assessment systems can also act as powerful business enablers: providing a baseline of risk compliance review immediately to individuals, immediately enabling low-risk activities with the proper controls, and triaging medium- and high-risk activities for more in-depth review from compliance professionals. While organizations vary widely in their adoption of automation within their risk management systems, some have developed end-to-end automated risk assessment tools that are fully traceable and designed to improve transparency between internal stakeholders. Importantly, these automated workflows introduce compliances signals earlier in product design cycles and enhance efficiencies within the organization.

Case Study

In highly regulated industries, risk classification can have pre-determined consequences. For example, in the financial industry, the distinction between low-, medium-, and high-risk classifications each require different actions and controls. To ensure automated risk assessment tools provide reliable outputs, a finance organization emphasizes that human-drafted inputs must be reliable. To support its workforce, it developed a control library that supports high-quality and consistent input prompts in the automated risk assessment tool. To emphasize human accountability, the agentic AI system is limited to a primary review, helping compliance professionals hone in on the risks that require more attention and safeguards.

For forward-thinking, accountable organizations, navigating and managing risk is no longer confined to legal and compliance teams—it's a strategic priority for growth, resilience, and trust that demands organization-wide ownership. Organizations want to bridge internal compliance programs and align risk management across various regulatory regimes, for example, across data protection, cybersecurity, trust and safety, and AI. Rather than completing risk assessments in isolation with multiple different teams and hand-offs, organizations increasingly want to integrate risk assessment considerations across various disciplines to simplify governance and develop integrated signals for product design. AI-powered systems can incorporate all relevant signals and details and draft various preliminary risk assessments for human review.

Finally, automating risk assessments can enable a more holistic consideration of both the risks and benefits of a product, application, project, or data processing activity. Given the significant benefits of AI technology, as well as its inherent risks, compliance teams must be able to weigh the potential loss of opportunities and benefits from not using AI against the potential risks of use to individuals or groups of individuals. Automation can also help ensure that risk assessment processes are able to be scaled and conducted consistently across different teams and functions with a varying range of technical understanding.

Considerations

- Designing AI-powered systems to support risk management processes requires high-quality, structured data reflecting historic internal risk management processes, considerations and decision-making. Many organizations do not have access to the amount of structured data and sources needed to design reliable AI-powered systems for risk management with high stages of automation. Organizations should carefully assess the data categories and assets needed to automate the identified elements of their risk assessment workflows.
- In addition to high-quality, structured data, organizations must build risk assessment methodologies, including risk scoring frameworks, to standardize risk analysis across various use cases and contexts. In addition, risk tolerance levels must be defined to inform the appropriate regulatory and technical considerations and mitigations. For example, organizations should have policies and procedures to address tensions and trade-offs between different risks and benefits such as tensions and benefits between privacy and trust and safety.
- Risk is not static. Organizations must account for how risk is constantly evolving and changing, including in AI-powered models and systems.
- Compliance professionals need to develop trust in automation tools, especially when tools assess risk or inform decision-making. Trust development in AI for compliance tools is an ongoing process. Project participants with mature tools discussed the value of starting with small, low-risk experiments and slowly expanding to larger projects with higher potential for risk.

4. Standardizing Policies, Controls, and Safeguards

Standard, manual compliance tasks rely on layers of human judgment and interpretation, which can be highly subjective and error prone. Automation can help translate organizational policies, governance controls, and safeguards into clear, repeatable processes that produce better outcomes and increase consistency in decision making. By embedding regulatory requirements, risk taxonomies, and control frameworks directly into workflows, AI-powered tools can support the uniform application of compliance controls across teams and use cases. This standardization not only improves efficiency, but also strengthens organizational accountability, as decisions become traceable, auditable, and grounded in predefined rules and methodologies.

In practice, automated systems allow organizations to define global controls and standards and apply them consistently in real time, facilitating faster, more reliable compliance assessments and enabling scalable oversight across complex operations. At the same time, codifying best practices and safeguards into structured systems enhances transparency and defensibility, positioning organizations to demonstrate to regulators that compliance is embedded into operations through clear, repeatable processes and governance mechanisms. However, standardization strengthens accountability only where the underlying rules, taxonomies, and methodologies are themselves sound, i.e., updated to reflect regulatory changes and appropriate to the context in which they are applied.

With proper guardrails and when developed and deployed responsibly, agentic AI systems can minimize the risk of human error, variability, and even misconduct by applying standardized logic across a variety of tasks. This can be particularly valuable in contexts that require consistent application of organizational or regulatory policies. For example, agentic AI can be deployed to automate due diligence requirements, such as know-your-customer (KYC) and anti-money laundering (AML) or monitor for instances of potential internal policy violations or regulatory noncompliance.

Considerations

- Organizations should manage model drift: where an AI model is trained on years of historic decision-making, updating it to reflect new regulatory or organizational policy developments requires careful management and fine-tuning to prevent the model from defaulting to outdated conventions.
- Standardization can propagate a flawed premise at scale. Organizations should ensure that the underlying rules, taxonomies, and methodologies are sound, and not treated as an after-thought.
- Much of the value of compliance judgment lies in the exercise of contextual discretion, particularly at the margins and in the case-by-case reasoning that regulatory requirements often anticipate, such as proportionality and balancing assessments. Organizations should standardize routine, well-defined tasks while preserving human discretion where the nature of the task genuinely requires it, consistent with a staged, risk-based approach to automation.

5. Enabling Real-time Audits

AI-powered systems can be designed to proactively produce audit trails, improving trust in the capabilities of the tool as well as trust in the organization from individuals, regulators, and other stakeholders. Automated workflows can also enable real-time auditing, which one project participant described as the 'holy grail of compliance'. Real-time auditing can bolster regulatory confidence and external audit readiness and provide organizations with the ability to anticipate, detect and correct anomalies, irregularities or control failures as they occur.

Considerations

- Despite the benefits that auditability drives, compliance professionals are concerned about the loss of confidentiality and the potential risk of exposure in regulatory investigations or litigation. This tension between improved documentation and the increased risk of discoverability (or undermining legal privilege) is a serious concern for compliance professionals. Some organizations respond to this tension by ensuring that certain high-risk processes are never completely automated. Regardless, organizations should create internal policies and procedures governing AI use to preserve legal privileges such as confidentiality when and where appropriate.

- Explainability is a non-negotiable requirement, especially for regulators and when decisions impact individuals. However, challenges remain with meaningful explainability when certain AI systems—particularly those using complex machine learning techniques such as agentic AI, are relied on for decision making and relevant outputs, as in some cases the relationship between a given input and the resulting output cannot be readily traced or reconstructed. The question then is what practical information, data input, model logic, chain of thought, and audit trails ought to suffice as meaningful explainability. Furthermore, as agentic AI systems become more interconnected and in circumstances where agents are interacting or collaborating with each other, it may become unclear which component is responsible for a particular decision or data use. Immutable logs of agents AI systems can provide organizations with detailed records for monitoring and analyzing agent behavior, ensuring human oversight, and enabling proper redress mechanisms. Log immutability also ensures that problematic agent trajectories and failures cannot be deleted, preserving them for analysis and compliance purposes.

6. Augmenting Human Expertise

AI-powered compliance tools are most powerful not as replacements for human judgment, but as force multipliers that elevate human expertise and enable informed decision-making. Throughout the project, a consistent theme emerged: AI excels at handling scale—analyzing vast datasets, surfacing relevant regulatory requirements, conducting preliminary risk assessments, and generating structured options—while freeing compliance professionals to focus on the nuanced, high-value work that requires judgment, context, and ethical reasoning. Rather than narrowing human involvement, AI shifts it upstream and downstream, enabling experts to engage more critically with outputs, interrogate risks earlier, and refine decisions with greater precision.

In practice, AI-generated insights, draft analyses, and decision-support tools provide a baseline of structured intelligence that helps humans identify blind spots, compare options, and challenge their own assumptions, ultimately leading to more deliberate and defensible compliance outcomes. Importantly, the "human as architect" model ensures that experts formulate the objectives, set the parameters, define risk tolerance, and validate outcomes, reinforcing accountability while leveraging AI to enhance, not dilute, professional judgment. By combining machine-scale analysis with human expertise, automation can create an evolved compliance cycle: better inputs leading to better outputs, which in turn enable more rigorous, consistent, and thoughtful decision making across increasingly complex regulatory landscapes.

Accountable organizations do not want to replace their compliance teams with automated workflows, they want to prioritize human capacity for the most high-value work and expand compliance coverage across the organization.

There was broad agreement among project participants that reliance on manual compliance processes means that not every project receives structured review. AI-powered tools can change that equation, making baseline compliance checks scalable and accessible in ways that weren't possible before. The pressures and opportunities for compliance teams to adopt automation mean that teams that fail to will be at a disadvantage in the near future.

Considerations

- Among project participants, a central risk is not that AI will replace human judgment, but that it will quietly reshape it—encouraging compliance professionals to confirm rather than critically evaluate automated outputs. AI tools can produce ostensibly confident and well-structured outputs, which can create a false sense of accuracy, causing users to accept recommendations without sufficiently questioning them. This is known as automation/ confirmation bias and can cause compliance professionals to overly defer to AI-powered systems instead of exercising independent judgment and verifying outputs.
- Some project participants raised concerns that over-reliance on automation in compliance functions can erode expertise, especially among less experienced compliance professionals. Younger or less experienced professionals may not develop the baseline expertise needed to challenge AI outputs effectively. Over time, compliance teams may lose the ability to detect subtle errors or edge cases.
- AI-powered systems must be designed with safeguards, ensuring humans are empowered and incentivized to validate AI outputs and challenge them when necessary rather than just accept them.

Considerations for Regulators and Policymakers

- **Automation has the potential to fundamentally reshape compliance from a reactive, event-driven function into a proactive, embedded governance capability.** By integrating compliance controls directly into operational workflows, AI enables continuous monitoring, early risk identification, and real-time interventions, reducing the need for retrospective enforcement while improving accountability. Regulators should recognize and support this shift, as it aligns with modern regulatory approaches that prioritize prevention, engagement, and improved outcomes over purely punitive models.
- **AI-driven compliance can strengthen demonstrable accountability and enable more effective, outcomes-based compliance.** Automating and enhancing compliance with AI can support a shift from seeing compliance as an administrative or process-driven burden to something that not only enables the business, but also more effectively achieves the underlying objectives of regulation. The risks associated with AI-driven compliance are highly context-dependent and should therefore be assessed through a proportionate, risk-based approach. Regulatory scrutiny should be calibrated to the level of risk and followed by practical measures to mitigate the risk, with particular attention given to systems that lack explainability or have the potential to materially impact individuals. For example, where full explainability cannot be achieved, are chain-of-thought reasoning controls sufficient for certain use cases and contexts? Taking such an approach enables organizations to realize the benefits of AI for compliance while focusing safeguards and oversight where they are most needed. Ultimately, effective use of AI for compliance and better compliance outcomes can enable sustainable digital trust, better performing products and services, and increased trust in innovative technologies.
- **Regulators play a critical role in shaping whether AI for compliance becomes a driver of accountability or a missed opportunity.** Rather than adopting overly cautious restrictions for the use of AI for compliance, policymakers and regulators should be more proactive in encouraging and emphasizing the responsible use of automation and AI tools to deliver effective accountability and compliance, which can ultimately provide better outcomes for people, customers, and society. Similar to incentives introduced for greater adoption and investment in PETs, AI for compliance solutions can be encouraged through clear, practical frameworks that enable and incentivize responsible experimentation and deployment. This can include regulatory guidance (e.g., the role of transparency and explainability, the extent of human-in-the-loop in design), regulatory sandboxes, and collaborative engagement with industry. This approach supports innovation while maintaining safeguards, ensuring that regulation reflects operational and technical realities and enables organizations to improve compliance outcomes in practice.

- **Regulators should embrace AI tools to maximize their effectiveness as demand for their time and expertise continues to grow.** Regulators will too have to become comfortable developing, deploying and monitoring AI tools to augment and where appropriate automate tasks and duties, such as complaints handling, investigations, information processing, providing AI-enabled guidance tools, and reporting on KPIs. This is particularly true as some regulators report increased public complaints bolstered by widespread public access to AI tools.[2]
- **Public sector organizations should also look to invest in AI for compliance tools and the necessary governance required for their responsible use.** Though this project's participants primarily consist of private sector organizations, the project's findings and best practices are equally relevant and applicable to public sector organizations and government bodies. Many such entities already process personal data to fulfill their statutory duties and are increasingly seeking to adopt AI tools to improve efficiency and outcomes. They are all also subject to privacy, cybersecurity, and other compliance requirements and should therefore invest in accountability programs and AI tools to help operationalize and deliver effective compliance.

2 Haeck, P., & O'Regan, E. (2026, April 10). EU complaints system buckles under pressure of AI. POLITICO; and Stan, A. M. (2026, May 29). Australia's workplace tribunal says AI-assisted claims have helped drive a 70% workload increase in three years. The Next Web.

Conclusion

The adoption of AI-powered tools represents a turning point for compliance functions, enabling them to evolve from constrained, reactive processes into scalable, proactive, and strategically integrated capabilities. By automating repetitive tasks, standardizing decision-making, and embedding controls directly into operational workflows, AI empowers compliance teams to expand their reach, improve outcomes, and enhance accountability. These tools not only strengthen core functions such as risk assessment, monitoring, and recordkeeping, but also generate real-time insights, continuous auditability, and improved transparency, reinforcing trust among regulators, stakeholders, and the public. At the same time, AI enables compliance professionals to refocus their expertise on high-value, judgment-intensive work, ensuring that human oversight remains central to decision-making.

These benefits, however, are not automatic. The very properties that make automation valuable (consistency, scale, and speed) are also what give risk to the risks and challenges this report has identified. Realizing the opportunity therefore depends how these tools are designed, deployed, and governed. This is also what makes compliance's relationship with AI distinctive. Compliance teams are not simply another function adopting new technology; they are the function responsible for governing AI across the organization, and increasingly for equipping the wider business to use it responsibly. Getting their own adoption right is therefore not only a matter of efficiency but of institutional accountability.

When deployed within strong governance frameworks and guided by principles of organizational accountability, AI for compliance becomes more than an efficiency tool. It becomes a foundational enabler of responsible data use, resilient governance, sustainable innovation in an increasingly complex regulatory environment, and a shared undertaking in which industry and regulators each have a role to play.

Appendix I: Use Cases

Below are use cases collected from CIPL members and workshop participants throughout the two workshops and individual interviews:

Regulatory Dashboards

Many organizations have developed internal dashboards to support their legal and compliance teams with streamlined and indexed regulatory and policy developments. Dashboards can be embedded with internal documents, allowing legal and compliance professionals to immediately assess the business impact of a new regulatory development or potential legislation. Additionally, the tool can suggest potential contractual, organizational, or technical solutions in response to regulatory developments. One organization designed its dashboard with human-review considerations—users can verify regulatory developments against linked sources and immediately review internal documents like contracts. This design helps create a culture within the organization that questions automation bias and reaffirms human professionals as accountable decision makers.

Document Review

A financial institution that is required to review high volumes of customer provided documentation designed an internal agentic AI system to reduce the manual time required to complete document review related tasks while improving the consistency of the analysis. The system is not designed to make any credit related decisions, it is only designed to support decision making by reviewing documents and generating internal compliance related documents based on internal guidelines and benchmarks. Despite having no decision-making function, the system is still designed to ensure human review, for example, human reviews verifies the accuracy of the system's document review.

Risk Assessments

An organization implemented a fully agentic, automated workflow replacing a previously manual risk review process involving multiple teams and handoffs. While the AI agent drives an automated risk assessment process, human input and review remains central to the management of the tool. The AI agent performs preliminary risk assessments, drafts preliminary DPIAs, and other assessments depending on the context, and presents these to human owners. The workflow can also advise on appropriate controls, access rights, and connected knowledge bases or data sources. Built on the concept of "humans as architects" with human-in-the-loop controls in place when material risk thresholds are met, smaller teams are generally required to manage the automated workflow.

Another organization has incorporated privacy, compliance, information security, and audit into one function. This team developed an automated, organization-wide risk assessment tool to evaluate privacy, security and compliance related risk. Users are guided to input the necessary information upfront to inform thorough analysis, and low-risk activities are provided with automatic approval. Medium-, high-, and material-risk activities are routed to gradually deeper level of human review, with material-risk activities requiring compliance committee review. The automated workflow is entirely traceable and provides transparency to all stakeholders. The organization has reduced assessment volumes by filtering out low-risk activities and simplified governance across the organization, allowing for accelerated decision making.

Data Subject Access Requests (DSARs)

Many organizations have replaced manual, email-based DSAR intake systems with automated systems to keep up with increasing volumes of DSARs in a timely and thorough manner. Automated and integrated systems can provide data subjects with immediate access to requested data in most circumstances and route more complex requests to compliance professionals for manual review.

Product Counseling

This tool helps drive faster input, earlier in the product development cycle by allowing product teams to input relevant product descriptions, context, and relevant facts and receive clear guidance that reflects the organization's governance and risk tolerance. The tool is trained to consider the various factors that inform privacy compliance considerations, including how a product works, how it is or is not different from what's been developed before, and what data is involved and where the data flows. While this AI tool allows product teams to get early feedback on compliance matters, the final product is subject to human review before deployment. The level and thoroughness of human review depends on the underlying assessment of risk involved.

Third-Party Vendor Risk Assessment Chatbot

One organization designed a chatbot that helps manage third-party vendor risk. Employees can explain what they what to do with the vendor and get access to immediate intelligence on the risks, as well as determine whether existing vendors are available for the task. The chatbot also supports with vendor contract management by identifying key provisions based on internal requirements. Finally, the integrated intelligence from the tool supports overall vendor risk management.

Internal Governance Chatbot

This chatbot allows employees to quickly query internal documents, policies, and procedures. One organization deployed an AI agent that supports employees with AI governance related matters. Employees can quickly see which AI tools are approved for certain uses cases, evaluate the Markdown file associated with an employee-built AI agent against organizational policies, or request permission for a new AI use.

Co-Worker Chatbot

Internal chatbots can also help employees with brainstorming, editing and drafting, and accurately completing compliance filings. These tools are trained on selected internal communications, for example, emails, messages, and internal memos, to learn advisory styles and judgment patterns. Project participants highlighted the distinction

between using these tools as editors, to refine and improve human capability, and as writers, where confirmation bias can cause over-reliance and reduced human capability.

Cookie Categorization Tool

AI-powered systems can review cookie-related details, assign them into categories based on purposes, suggest the necessary controls, and maintain auditable records. These systems can also be programed to scan websites for cookies in real-time to ensure thorough regulatory compliance.

Sales Call Compliance Tool

This tool monitors and ensures that employee sales calls comply with regulatory requirements, requiring that certain disclosures are made and preventing misrepresentations.

Due Diligence and Security Questionnaires

Agentic AI systems can automate responses related to due diligence and security compliance questionnaires, which usually contain hundreds of repetitive questions and take several days to complete manually. Controls are carefully designed to reduce the chances of inaccurate outputs while citing relevant documents for outputs in an evidence folder.

Cybersecurity Supply Chain Monitoring

Agentic AI systems can automate time-consuming and ever-evolving cybersecurity compliance tasks. One organization deployed an agentic AI system that monitors subcontractors and maintains up-to-date documentation, saving its legal team hundreds of hours. The system is developed with agent-as-a-judge and RAG frameworks to improve response accuracy and establish evidence measures for its outputs. One project participant highlighted that the output of agentic AI systems is only as resilient as the sources it is provided with.

Appendix II: Key Considerations for Designing and Deploying Automated Compliance Tools

1. Adopt a staged, risk-based approach

- Use **maturity models (e.g., 5 stages of automation)** to gradually scale AI adoption.
- Align level of automation with:
  - Priorities and impact on the organization
  - Risk level
  - Use case complexity
  - Regulatory constraints

2. Define meaningful human oversight

- Move beyond generic human-in-the-loop to **"the right human in the loop"** with:
  - Defined expertise
  - Clear accountability
  - Escalation pathways
- Consider the **"right loop"**, focusing on when human review is required, such as when material risk thresholds are met or when individual rights are impacted.

3. Build strong data foundations

- Ensure:
  - High-quality, structured data
  - Access to historical compliance documentation
- Recognize that **data quality and sources directly determine AI performance**

4. Prioritize explainability and auditability

- Design systems that:
  - Produce **defensible audit trails**
  - Allow investigation of outputs and sources
  - Support scrutiny from individuals, customers, regulators and other stakeholders

5. Invest in AI literacy and workforce enablement

- AI literacy is **critical to adoption and responsible use:**
  - Training programs
  - Embedding technical expertise into compliance teams

6. Align AI with organizational accountability frameworks

- AI-powered compliance tools should be designed to enhance organizational accountability by:
  - Supporting risk assessments
  - Enabling transparency
  - Strengthening governance processes

7. Carefully select and prioritize use cases

- Focus on:
  - **Low-risk, high-volume tasks first**
  - Slowly evolve to high-value use cases with clear business value
- Maintain human oversight for:
  - High-risk decisions
  - Sensitive or impactful outcomes

8. Embed continuous monitoring and improvement

- AI systems require:
  - Ongoing testing and validation
  - Model tuning and refinement
  - Real-time monitoring

9. Manage legal and evidentiary implications

- Be aware of:
  - Documentation risks (e.g., privilege erosion, discoverability)
  - Liability concerns tied to AI-powered decision-making
- Design processes to **balance transparency with legal risk**

About the Centre for Information Policy Leadership

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, 85+ 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 for 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. Nothing in this document should be construed as representing the views of any individual CIPL member company or Hunton. This document is not designed to be and should not be taken as legal advice. For more information, please see CIPL's website at:

http://www.informationpolicycentre.com/



DC Office
2200 Pennsylvania Avenue
Washington, DC 20037
+1 202 955 1563

London Office
30 St Mary Axe
London EC3A 8EP
+44 20 7220 5700

Brussels Office
Avenue des Arts 47-49
1000 Brussels
+32 2 643 58 00

© 2026 Centre for Information Policy Leadership. All Rights Reserved.

26519.1126.R2