ISSUE N.15 | DECEMBER 2025

Diritto intelligente

DLA PIPER

Contents

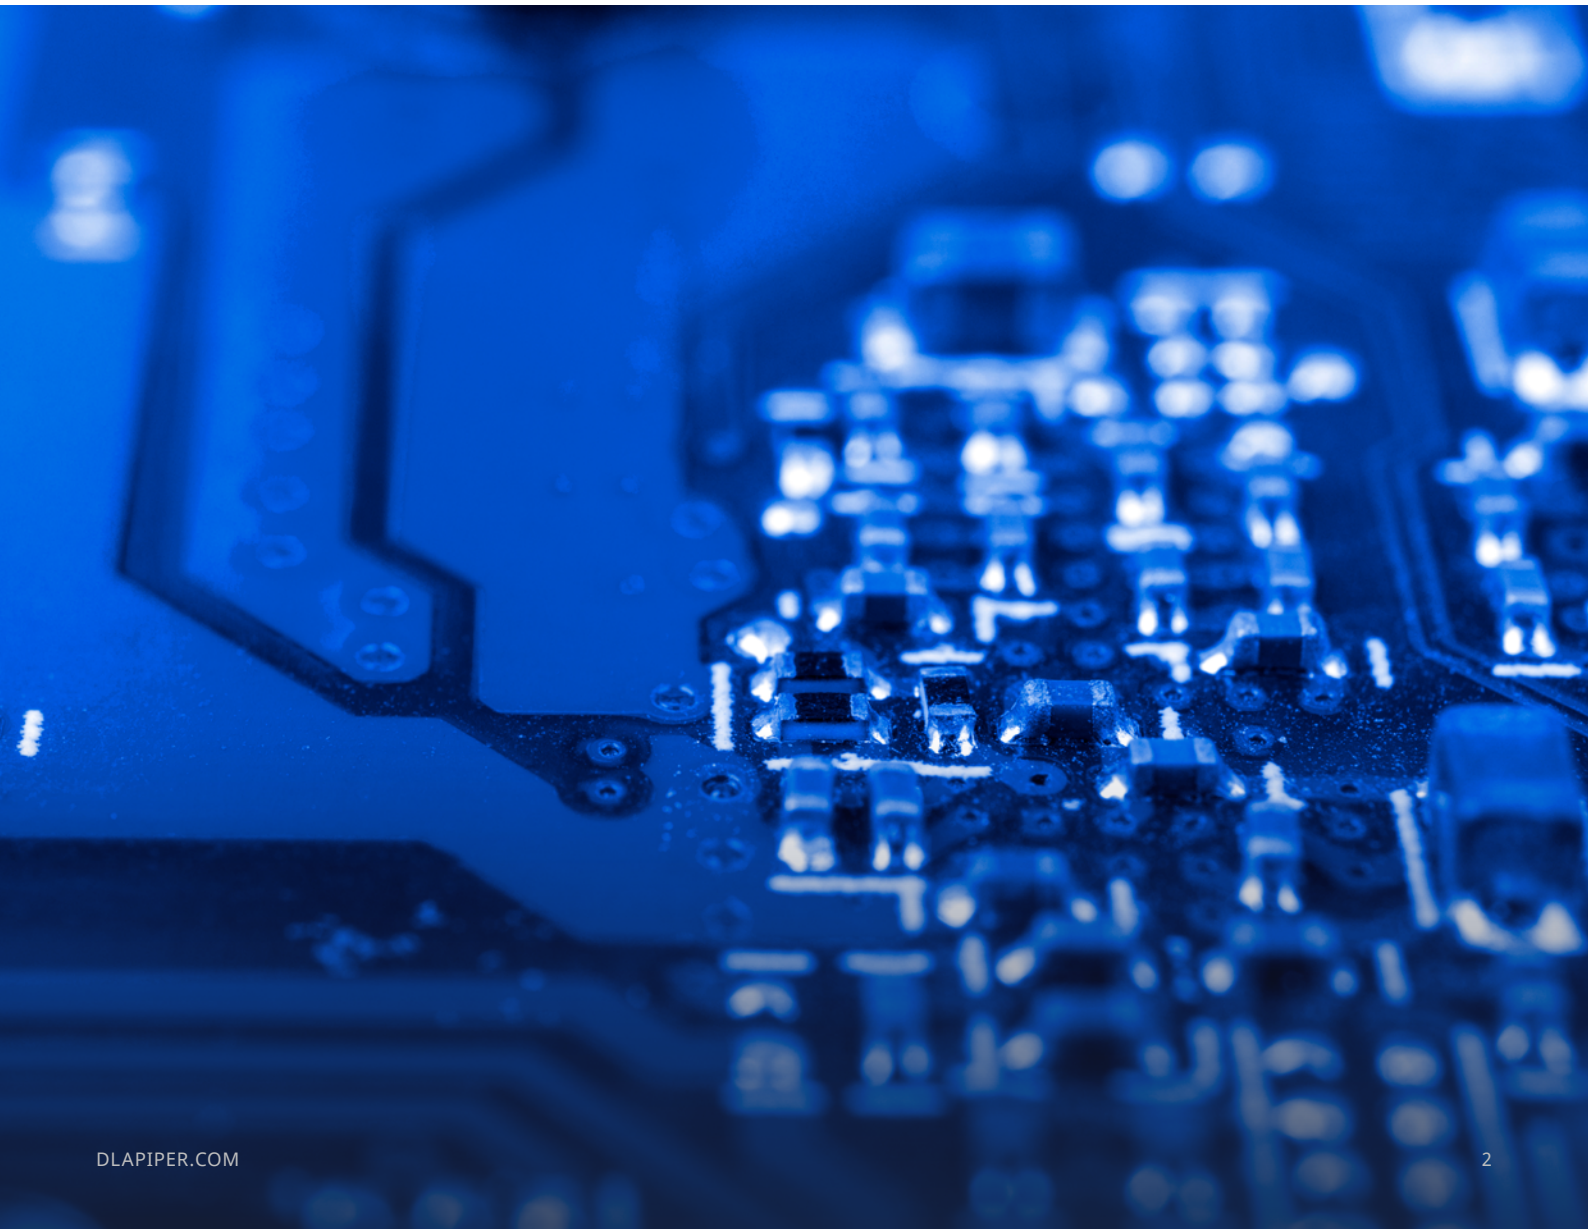

Editorial

December is traditionally a time to pause, look back, and look ahead. This issue of *Diritto Intelligente* does all three, offering a snapshot of how European digital law is moving from rule-making to execution, just as we enter what will likely be the first fully operational year of AI governance.

At the policy level, the Digital Omnibus proposal on AI training based on legitimate interest signals a pragmatic shift. It acknowledges that innovation needs legal certainty, while still anchoring AI development within GDPR safeguards. At the same time, the broader Package on Simplification shows a clear political direction: fewer overlapping rules, clearer interfaces between GDPR, the AI Act, data and cybersecurity laws, and a stronger focus on efficiency without lowering protection standards.

This transition from theory to practice is echoed in the EDPS guidelines on AI risk management. Here, compliance is framed as a continuous governance process, covering explainability, bias, security, data subject rights and vendor accountability. It confirms that AI risk management is no longer a technical niche, but a board-level responsibility.

The same practical approach emerges in sector-specific contributions. EUIPO's Early TM Screening illustrates how AI can concretely support professionals without replacing legal judgement. The Italian Council of State decision on AI in public procurement clarifies that innovation is welcome, but only under informed authorization, transparency and human oversight. In insurance, AI appears both as a powerful anti-fraud tool and as a new source of risk, requiring careful alignment with the AI Act and supervisory expectations.

Taken together, these contributions tell a coherent story. Europe is not slowing down AI. It is asking for maturity, responsibility and proof. As holiday season and the new year approach, the message is clear: 2026 will be less about announcing AI strategies and more about showing that they work, legally, ethically and operationally.

Warm wishes for the holiday season and for a new year where innovation and trust grow together.

Giulio Coraggio
Location Head of the Italian Intellectual Property and Technology Department at DLA Piper

AI training based on legitimate interest: Is the digital omnibus proposal enough?

Author: *Giulio Coraggio*

The European Commission's latest Digital Omnibus package introduces a significant and much-debated idea: allowing **AI training based on legitimate interest**, under Article 6(1)(f) GDPR, accompanied by a new Article 88c. The proposal formalises something many expected – that training AI systems or AI models on personal data *may* rely on legitimate interest as a legal basis.

But the key question emerges immediately: **is this enough?**

And does the wording of the proposal offer the clarity that the European AI ecosystem needs?

At first sight, the Commission's move appears to acknowledge practical needs. Innovation in AI, especially large language models and foundation models, depends on vast amounts of data. Given the tension between AI development and strict consent-based frameworks, the introduction of a dedicated legal pathway under legitimate interest seems both necessary and overdue.

Yet, when we examine the proposed wording closely, several uncertainties remain – and these uncertainties may determine whether the provision becomes a meaningful tool or a regulatory puzzle.

A clear step forward, but framed in cautious terms

The draft states that where the processing of personal data is necessary for the development or operation of an AI system or AI model, such processing "**may be pursued for legitimate interests**" under Article 6(1)(f) GDPR.

The choice of the word *may* is not neutral. It does not create an unconditional presumption.

It does not replicate the clarity and simplicity of the ePrivacy Directive's "soft opt-in," which gives controllers a clean, objective exception.

Instead, it maintains the existing structure of legitimate interest assessments and leaves several determinations – necessity, proportionality, balancing of interests – to the controller *and*, ultimately, to data protection authorities.

In other words: **this is not a safe harbour**.

Conditions that may create new layers of interpretation

Article 88c introduces conditions that appear reasonable in principle – organisational and technical safeguards, minimisation, transparency, and non-disclosure of residually retained data.

However, the practical implementation of these requirements raises several challenges:

- How detailed must minimisation be during dataset selection?
- To what extent must controllers ensure "enhanced transparency" to data subjects, especially if training relies on publicly available content?
- How should the "unconditional right to object" be operationalised for model training?
- Where is the line between residual retention and functional necessity for model integrity?

Each of these questions invites interpretation. And interpretation invites fragmentation.

For companies developing or deploying AI systems at scale, the risk of divergent national assessments is real – and this undermines the very goal of harmonisation that the Digital Omnibus seeks to achieve.

Could the EU Commission have been braver?

The industry's expectation was that the Digital Omnibus would go further. Many hoped for a straightforward provision establishing that AI training *is presumed* to rely on legitimate interest, provided that safeguards are met – much like the ePrivacy soft opt-in.

A stronger presumption would have:

- reduced DPAs' discretion,
- offered developers clearer legal certainty, and
- prevented inconsistent interpretations across Member States.

Instead, the Commission opted for incrementalism rather than boldness.

Whether this choice reflects political caution, ongoing litigation, or a desire to preserve regulatory flexibility remains open to debate.

Impact on ongoing disputes between LLM providers and DPAs

A central question now is whether Article 88c, as drafted, will meaningfully affect ongoing disputes between LLM developers and European regulators.

On the one hand, it recognises legitimate interest as a viable legal basis for training – which several DPAs have challenged.

On the other hand, the safeguards and discretionary elements built into the provision give regulators wide room to scrutinise implementation.

This means that disputes may shift from "**Is legitimate interest allowed?" to "Does your implementation satisfy the criteria?**" – a more nuanced, but not necessarily easier, debate.

Final considerations

The proposal for **Digital Omnibus legitimate interest AI training** is undeniably a step in the right direction. It offers a structured legal basis for a practice that is essential to the development of European AI capabilities.

But the provision is cautious.

It avoids clear presumption.

It preserves DPA discretion.

And it opens the door to interpretive variability.

Whether this is the best balance the EU could strike – or simply the most politically feasible option – remains an open question.

What is clear is that the conversation is only beginning.

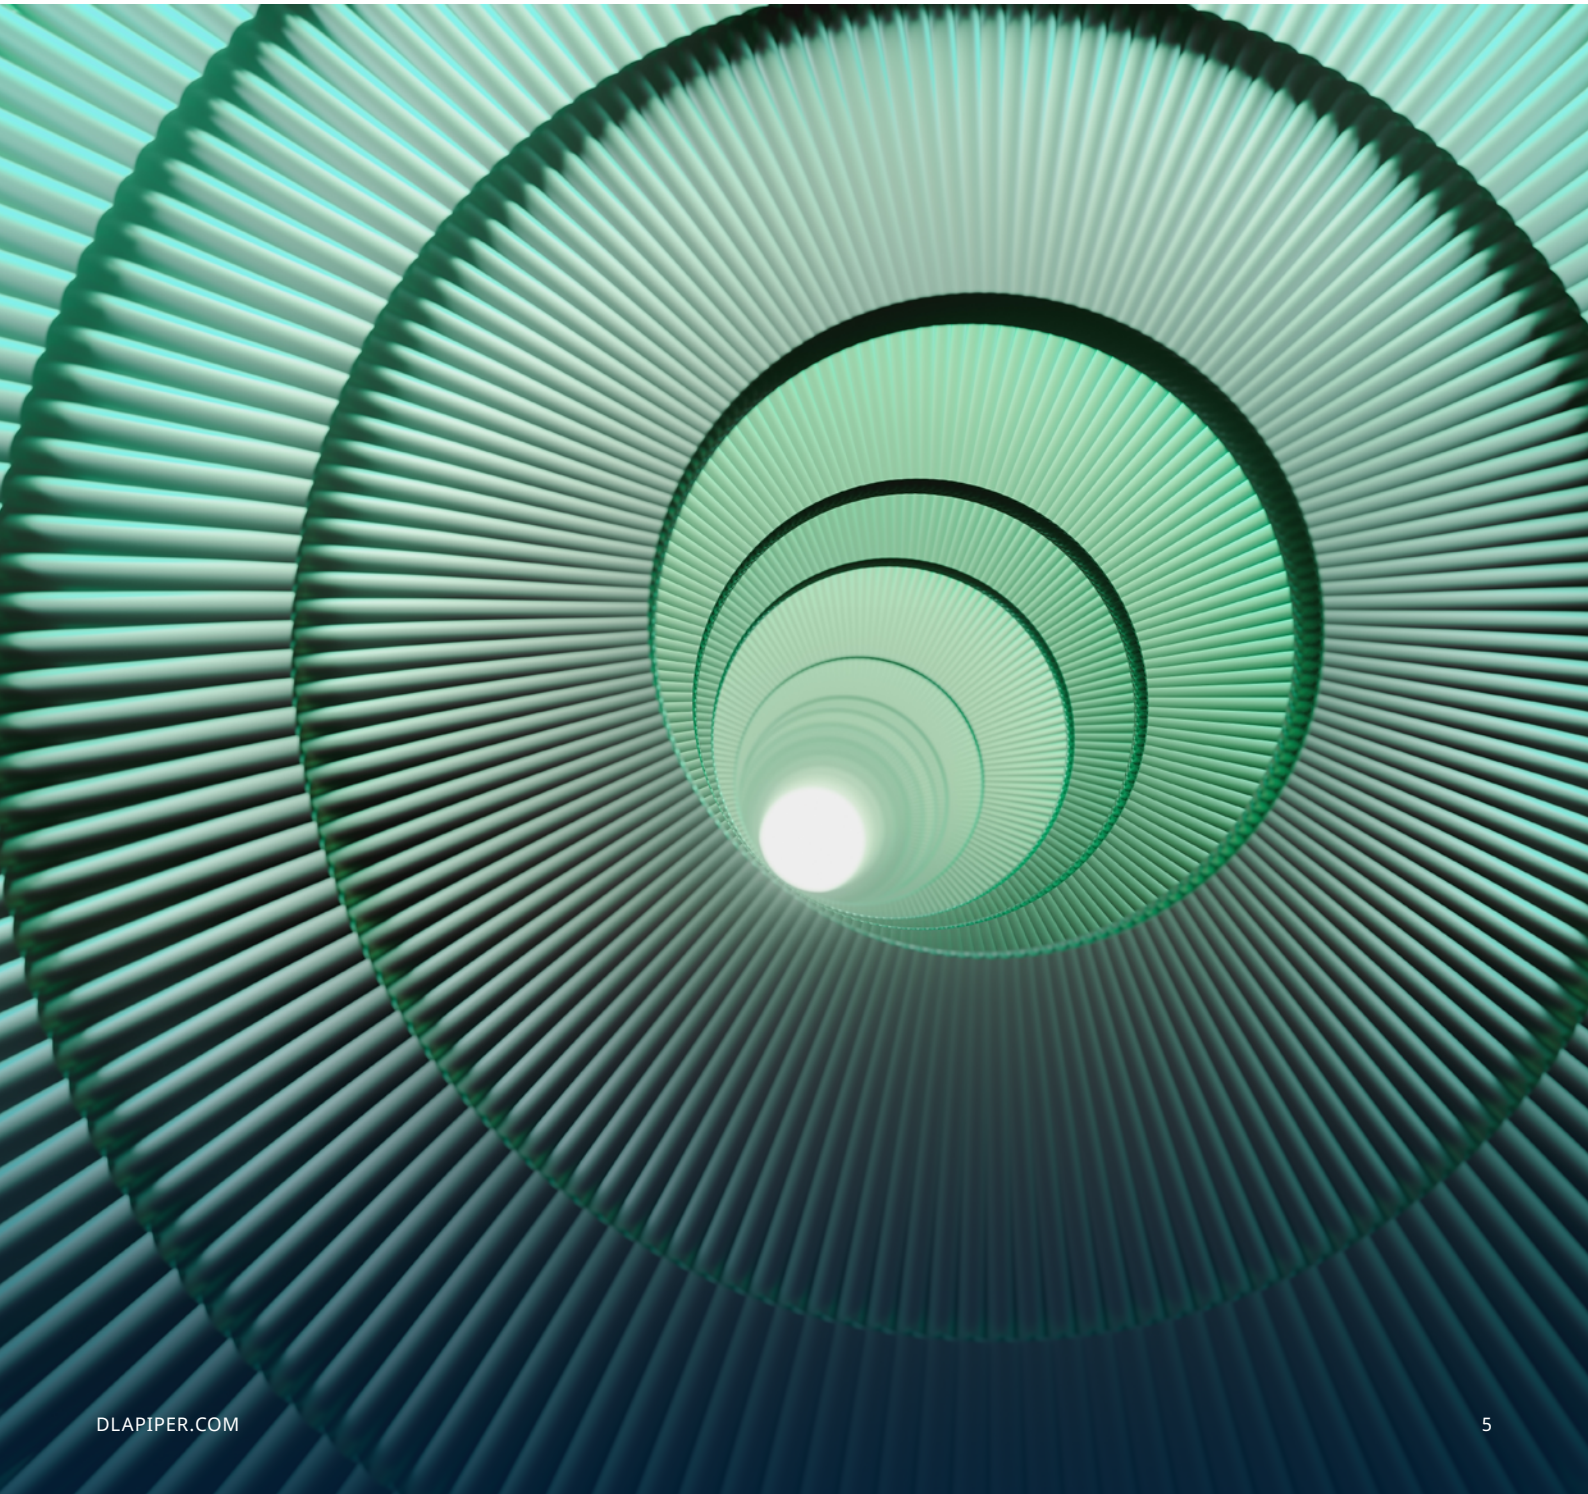

The EU's digital package on simplification: Streamlining GDPR, AI and Data rules

Author: *Giulio Coraggio*

The Digital Package on Simplification, proposed by the European Commission, updates key EU laws – including the GDPR, the AI Act, the Data Act, the NIS2 Directive, and the ePrivacy Directive – to modernize and simplify the entire European digital regulatory framework.

Also known as the Digital Omnibus, this proposal aims to streamline compliance, remove overlapping obligations, and reduce administrative burdens, while preserving the EU's high standards for data protection and digital trust.

By consolidating outdated rules and harmonizing legal obligations, the Digital Package on Simplification seeks to create a more efficient, innovation-friendly digital environment that benefits both businesses and regulators.

A political push for a "Simpler and Faster Europe"

The Digital Package on Simplification is part of the European Commission's broader political agenda titled "*A Simpler and Faster Europe*."

This initiative follows the Draghi and Letta reports on European competitiveness, which stressed that excessive regulatory layering undermines innovation and growth.

Responding to repeated Council conclusions in 2025, the Commission pledged to rationalize the EU's digital acquis by merging or repealing redundant texts and clarifying the interaction among existing digital regulations. The Digital Omnibus is the first concrete result of this simplification drive.

What the digital package on simplification proposes

According to an unofficial version of the document leaked from the EU Commission, the proposal introduces a single omnibus regulation that consolidates several major instruments.

It amends:

- GDPR (Regulation 2016/679)
- AI Act (Regulation 2024/1689)
- Data Act (Regulation 2023/2854)
- NIS2 Directive (Directive 2022/2555)
- ePrivacy Directive (Directive 2002/58/EC)

It repeals:

- Platform-to-Business (P2B) Regulation
- Data Governance Act (DGA)
- Free Flow of Non-Personal Data Regulation
- Open Data Directive
- By merging these acts into a unified structure, the Digital Package on Simplification eliminates inconsistencies, simplifies reporting requirements, and harmonizes definitions across the digital legislative corpus.

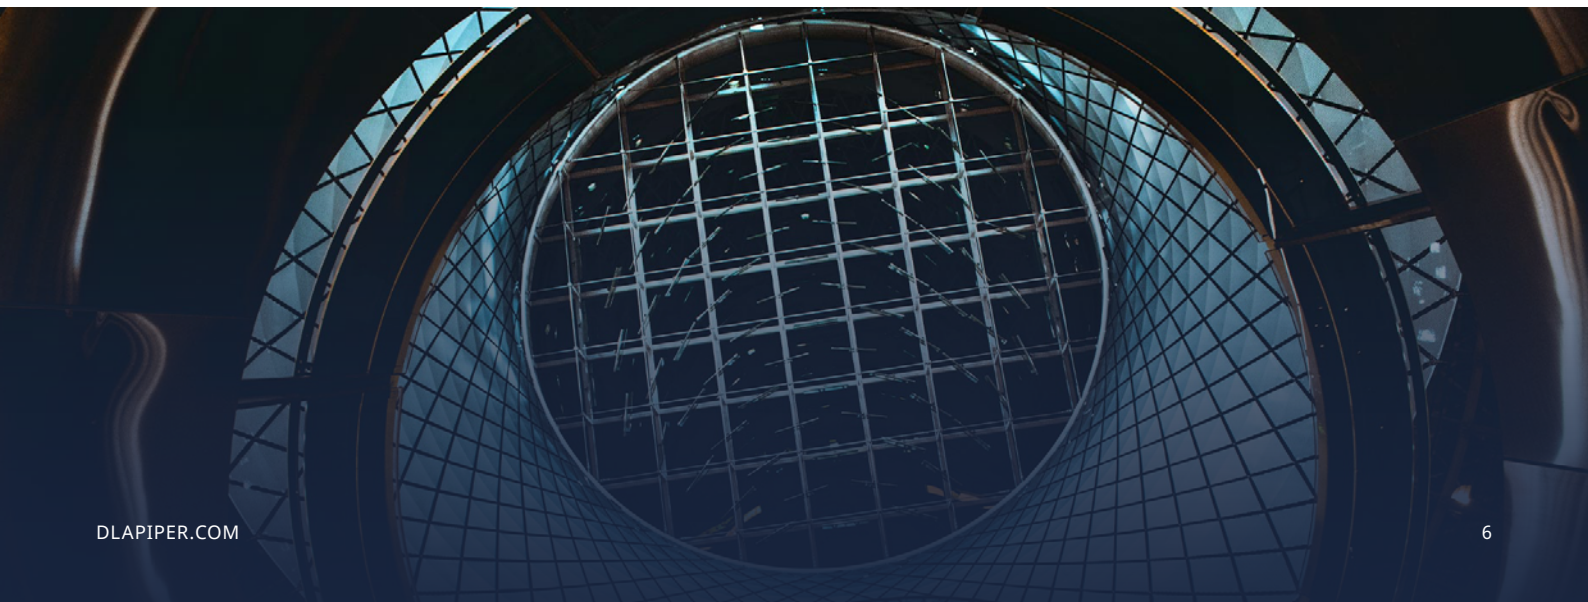

The three pillars of Simplification

The Digital Package on Simplification is built around three strategic pillars:

- Data framework consolidation
- Unified incident reporting
- Alignment of AI and data protection rules

1. Streamlining the data framework

The proposal merges the Data Governance Act, the Open Data Directive, and the Free Flow of Non-Personal Data Regulation into the Data Act, establishing a single point of reference for data sharing and reuse.

Key innovations include:

- Turning the registration system for data intermediaries into a voluntary trust framework within the Data Act.
- Consolidating rules for data altruism and public-sector data reuse.
- Clarifying cloud-switching and interoperability provisions.

This simplification is expected to cut administrative costs and improve legal predictability for companies working with data across borders.

2. One-stop system for incident and breach reporting

One of the most tangible improvements in the Digital Package on Simplification is the creation of a single EU- wide platform for incident and data-breach notifications, to be managed by ENISA.

This "*report once, share with all*" mechanism enables companies to fulfil obligations under the GDPR, NIS2, DORA, and Digital Identity Regulation simultaneously.

It will drastically reduce duplicate reporting, ease coordination between authorities, and increase efficiency – all while maintaining existing legal competences.

3. Aligning AI compliance with data protection

To reconcile the AI Act and GDPR, the proposal introduces clarifications on:

- The concepts of personal data and pseudonymization;
- The lawful use of personal data for AI training, under legitimate safeguards;
- Streamlined obligations for low-risk data processing.

These clarifications respond to long-standing industry concerns about legal uncertainty around AI training datasets and ensure a balance between innovation and privacy.

The end of the Platform-to-Business Regulation

The Platform-to-Business Regulation will be repealed, as its objectives are now fully achieved by the Digital Markets Act (DMA) and the Digital Services Act (DSA). This repeal reduces duplication and brings all platform governance rules under the same digital policy framework, improving consistency and enforcement.

Economic impact: Reducing red tape

The Digital Package on Simplification is also an economic reform.

According to the Commission, it will generate:

- EUR1 billion in annual savings;
- EUR1 billion in one-off savings;
- EUR4 billion total savings by 2029.

Small and medium-sized enterprises (SMEs) and small mid-caps will benefit most from reduced compliance obligations and simplified reporting mechanisms, strengthening Europe's digital competitiveness.

Legal basis and fundamental rights

Grounded in Articles 114 and 16 of the TFEU, the Digital Package on Simplification safeguards both market integration and privacy protection.

The Commission stresses that the initiative does not weaken the GDPR or the EU Charter of Fundamental Rights, but instead ensures a clearer and more coherent application of existing standards across all digital regulations.

A new chapter: Toward a Digital Acquis 2.0

The Digital Package on Simplification marks a fundamental shift from regulatory expansion to consolidation.

It sets the foundations for a Digital Acquis 2.0 – a unified, transparent, and innovation-oriented legal framework that strengthens Europe's position as a global leader in digital governance.

If adopted, the Digital Package on Simplification could become a model for how the EU modernizes complex legislation without compromising its core values of privacy, security, and accountability.

AI risk management: Seven key data protection lessons from the EDPS guidelines

Author: *Marianna Riedo*

As companies rapidly adopt artificial intelligence ("**AI**"), the need to manage its significant data protection risks has become a critical boardroom issue. On November 11, 2025, the European Data Protection Supervisor ("**EDPS**") published its "Guidance for Risk Management of Artificial Intelligence systems" ("**Guidance**"), to provide European Union Institutions, Bodies, Offices and Agencies ("**EUIs**") with some practical advice on how to ensure compliance when developing or using AI systems.

While this guidance is officially directed at EUIs, its principles and frameworks (which are build on the ISO 31000:2018 risk management methodology), are highly relevant for the private sector. The Guidance serves as a best-practice blueprint for any company looking to build a robust and legally compliant AI governance strategy.

1. Making interpretability and explainability the AI system's foundation

The EDPS considers interpretability and explainability not as interchangeable buzzwords, but as distinct, crucial concepts for understanding and trusting AI. As defined in the guidance, **interpretability** is the degree to which a human can comprehend how a model works and grasp the connections between its inputs and outputs. **Explainability**, on the other hand, is the ability to clarify *why* a model made a specific decision in a way that is accessible to an end-user.

A system that is interpretable and explainable allows AI providers to build confidence with customers, demonstrate compliance to regulators, enable effective audits, and more easily correct errors. This means that AI systems require a built-in solution to query *why* a decision was made, not just generate a clear and straight forward output.

2. Tackling bias from every angle: data, algorithms, and interpretation

The EDPS clarifies that the principle of fairness requires that personal data not be processed in a way that is "unjustifiably detrimental, unlawfully discriminatory, unexpected or misleading to the data subject." It makes clear that bias is not a monolithic problem but a multi-faceted risk that can emerge at different stages of the AI lifecycle. The document identifies five distinct root causes of bias:

- **Lack of Data Quality:** AI systems operate on a "garbage in, garbage out" principle. Inaccurate, incomplete, or poorly labelled training data can lead to flawed and biased outputs.
- **Bias in Training Data:** This can arise from historical biases reflected in the data or sampling errors that result in datasets that are not representative of the real-world population.
- **Overfitting:** This technical risk occurs when a model learns the training data, including its noise and outliers, so perfectly that it cannot generalize to new data, leading to poor and often biased real-world performance.
- **Algorithmic Bias:** The very design of the AI system, including the choice of mathematical functions or algorithms, can be inherently biased and produce unfair results, regardless of the data quality.
- **Interpretation Bias:** Even with a perfect model, human analysts can introduce bias by drawing skewed or incorrect conclusions from the AI's outputs, often influenced by their own preconceptions.

Operationally, this requires a cross-functional "responsible AI" working group that includes data scientists, legal counsel, and business line owners, with the mandate to challenge assumptions at each stage of the AI lifecycle.

3. Understanding the two faces of accuracy

The EDPS guidance highlights a critical distinction that is often lost in translation between legal and technical teams: the difference between legal **accuracy** and **statistical accuracy**. Under data protection law, accuracy means personal data must be factually correct, while statistical accuracy is a performance metric measuring how often a model's predictions are correct. This distinction is vital. An AI model can have very high *statistical* accuracy but still produce legally *inaccurate* personal data, creating a serious compliance gap. This is especially true for generative AI tools. Businesses must implement verification measures, such as human oversight, to ensure the factual accuracy of AI outputs that constitute personal data.

4. Navigating the tension between data minimisation and AI's data hunger

There is an inherent conflict between the core data protection principle of data minimisation and the fact that most AI systems require vast datasets to learn effectively. To address this, the EDPS suggests several technical mitigation measures, including data sampling (using a representative subset of data), anonymisation, pseudonymization, or synthetic data to reduce the amount of identifiable personal data processed.

5. Updating security frameworks for AI- specific threats

The guidance correctly notes that AI systems introduce unique security vulnerabilities that go beyond traditional IT security threats, including:

- **Training Data Disclosure:** Attackers can use techniques like model inversion and membership inference to query a model's output and reverse-engineer it to reveal sensitive personal data that was part of the original training set.
- **Data and Model Poisoning:** A malicious actor could intentionally manipulate the training data or the model itself to introduce hidden biases, backdoors, or critical errors.
- **API Vulnerabilities:** Poorly secured Application Programming Interfaces (APIs) that provide access to the AI model can become a major vector for personal data leakage.

Corporate cybersecurity playbooks must be expanded to account for these new attack vectors, as securing the AI model, its training data, and its APIs is now just as critical as protecting traditional databases and networks.

6. Operationalizing data subject rights for the AI Era

AI systems pose significant technical challenges to fulfilling fundamental data subject rights. The core difficulty, as the Guidance points out, is how to identify and then erase or correct specific personal data once it has been absorbed in the parameters of the model. The EDPS introduces "machine un-learning" as a potential technical solution, to develop methods that selectively make a model "forget" specific data points it was trained on without having to retrain the entire model from scratch. When machine un-learning is not viable, output filtering can be used, to detect and block personal information before reaching users though real-time scanning. This means that an erasure request may now trigger a complex and potentially costly model retraining or filtering process, creating a significant governance gap if not planned for in advance.

7. Demanding transparency when procuring AI systems

For the many companies that procure third-party AI systems rather than building them in-house, the EDPS guidance serves as a powerful checklist for due diligence. Simply trusting a vendor's claims is not enough. Rather, the AI vendor should be required to provide:

- General documentation on what the AI system does and how its underlying algorithms work.
- Specific information on how the system addresses transparency and explainability.
- Documentation on cybersecurity measures related to model integrity.
- Details on the provider's data governance practices, including how the training data was sourced and processed.
- The results from validation and testing procedures, including performance metrics on fairness and bias across different demographic groups.

Conclusion: From risk management to responsible innovation

The central message from the EDPS is clear: deploying AI requires a shift from a defensive, compliance-focused posture to a proactive and systematic risk management culture. This is not about stifling innovation with red tape; it's about enabling sustainable and trustworthy innovation by managing its inherent risks from a position of strategic foresight.

Early TM Screening: a new AI ally for professionals and enterprises

Author: *Noemi Canova*

EUIPO has recently introduced a new tool that has the potential to become a major aid for those working in the field of trademarks. This is Early TM Screening, launched as part of the SP2030 strategic plan, whose main objective is to make EU trademark registration process simpler, more transparent and reduce the possibility of mistakes. Bringing this goal closer is the use of artificial intelligence: the platform, in fact, relies on AI to carry out a preliminary diagnosis of the main issues that a trademark may encounter during the examination phase.

Accessing the screening is easy: you simply enter the sign to be analyzed and select the relevant classes of goods or services. The system then performs an automatic assessment that includes both the research for potential conflicts with earlier rights (EU trademarks, national trademarks, domain names, company names) and an evaluation of possible absolute grounds for refusal, such as descriptiveness, lack of distinctiveness, deceptiveness, or conflict with public policy or accepted principles of morality, as well as checks for any interference with specific rights such as geographical indications. This entire phase is accompanied by the possibility of downloading a summary report and seamlessly proceeding to the actual filing form.

Compared to the checks already available in online filing modules, Early TM Screening stands out for its ability to bring together in a single dedicated space a series of verifications that - until now - were scattered across different sections or required external tools. In particular, the plug- in based on the TMview tool allows for a quick and up-to-date comparative assessment, while other AI modules help identify conceptual similarities and automatically compare the sign with previous EUIPO's decisions in similar cases. The result is a richer and more predictive solution that offers users more landmarks than in the past. Naturally, **the platform does not aim to replace the interpretative and strategic work of experts**: the results are purely informative and not exhaustive, and the absence of reported issues does not equate to a guarantee of registrability. **Algorithmic analysis, although sophisticated, cannot capture every subtlety of potential conflicts, finer conceptual similarities or peculiarities of the relevant markets. For this reason, the technology is designed as a preliminary phase**, useful for refining initial choices and anticipating possible obstacles, but **not sufficient to replace a full clearance search or an in-depth legal assessment**.

Equally noteworthy is the new landing page dedicated to the innovative screening system, designed to provide users with educational support: in addition to tutorials and illustrated infographics, it features concrete examples of cases in which a trademark may be considered descriptive or non-distinctive, suggestions on possible solutions, and a structured list of the most frequent absolute grounds for refusal. This initiative clearly reflects the EUIPO's intention to guide users and professionals towards a more informed and structured approach, reducing the number of applications requiring subsequent corrections or clarifications.

In conclusion, Early TM Screening is not just a new digital tool, but a significant step in the modernisation of trademark protection in Europe. The automated pre-evaluation makes it possible to identify potential issues early on, facilitating a more solid and informed preparation phase. At the same time, it is a support that **must be combined with legal expertise, not replace it: only the combination of algorithmic analysis and professional judgment can ensure a truly effective filing strategy**. For law firms, this tool presents a valuable opportunity: on the one hand, it allows them to offer clients immediate and accessible first screening; on the other, it enables them to establish from the outset a clearer and more transparent dialogue on risks, opportunities and operational choices. A synergy between technology and consultancy which, if properly harnessed, can contribute to making the entire EU trademark ecosystem more efficient, more predictable and more responsive to the concrete needs of the market.

Artificial intelligence in public procurement: Insights from the Italian Supreme Court Decision

Author: *Dorina Simaku*

The Italian Administrative Supreme Court, the Council of State, issued the decision No. 8092/2025 which offers an important contribution to the ongoing discussion on the use of artificial intelligence ("**AI**") in public procurement and on the role of public authorities in authorizing and supervising such technologies. While the court ultimately rejected the appellant's claims, it took the opportunity to outline a set of principles that contracting authorities should follow when an economic operator proposes to use automated systems in preparing its bid or performing the contract.

1. The case: AI in the technical offer and the commission's discretion

The case arose from a tender for cleaning and sanitation services for entities within the National Health Service. The successful bidder stated that it intended to use AI tools as part of its operational model. The appellant argued that the evaluation committee had given disproportionate weight to this element and questioned the reliability of the assessment. The Council of State clarified that the committee had based its evaluation on a range of technical factors, and that the use of AI, while innovative, had not been decisive.

This point is significant: technological innovation cannot be assumed to be inherently advantageous, nor should it automatically raise concerns. It is for the contracting authority to determine whether the proposed technology aligns with the contract's objectives and complies with the legal obligations binding the administration.

2. The legal framework: authorization, transparency, and human oversight

The use of AI in public procurement is governed by a layered regulatory framework, starting with the Legislative Decree 36/2023 ("**Public Contracts Code**" or the "**Code**") and extending to the more recent Law No. 132/2025 on AI in public administration.

Article 30 of the Code expressly allows public authorities to rely on automated solutions, including AI systems, provided specific principles are respected. Automated decision-making must be transparent and understandable for both the authority and the economic operators involved. AI systems cannot be used to make fully automated decisions: human oversight remains mandatory. Contracting authorities must also ensure that the logic and functioning of the systems are knowable, providing access to source code and technical documentation, at least to the extent compatible with intellectual property protections.

Law No. 132/2025 further strengthens this framework by stating that AI may support—but never replace—administrative decision-making. A natural person must always remain responsible for the decision. As a result, adopting AI technologies in procurement requires a prior assessment by the contracting authority, which must authorize their use only after understanding their logic, functionality, and operational implications.

3. The role of the authority's consent and authorization

Whenever the use of AI involves the processing of personal data – as in the healthcare services at issue – the regulatory requirements become more stringent. Article 35(5-bis) of the Code requires economic operators to submit their consent to data processing through the virtual dossier system. However, this alone is not enough when sensitive data are involved.

In such cases, it is the public authority that must expressly authorize the use of AI, after verifying compliance with the principles of technological neutrality, transparency, cybersecurity, and data protection set out in Article 19 of the Code.

Authorization therefore plays a central role. It is not a mere formality: it is the moment in which the administration fulfils its duty to safeguard the public interest, assessing the level of control it will be able to maintain over the AI system during contract performance.

4. Algorithmic transparency and control over automated processes

A closely related issue is algorithmic transparency. Article 30 of the Code requires authorities to ensure that the logic behind automated systems used in procurement procedures is understandable. This means that a contracting authority cannot simply rely on a bidder's general description of the technology. It must have access to the technical information necessary to evaluate how the system works and how it will be supervised.

Case law on access to administrative records – such as TAR Lazio Decision No. 11335/2018 – makes it clear that generic claims of technical confidentiality cannot restrict the right of access, especially when transparency is essential to verify the legality of administrative action. This implies that AI solutions must be subject to real and effective scrutiny, compatible with industrial property rights but sufficient to ensure full administrative accountability.

5. Data protection and the authority's responsibility

Using AI in public procurement inevitably involves the processing of personal data. Law No. 132/2025 requires such processing to be lawful, fair, and transparent, and mandates that data subjects receive clear and understandable information. In sectors involving highly sensitive data, such as healthcare, the authority must also ensure that the operator uses AI systems strictly aligned with the purposes for which the data were collected, and that any risks are properly mitigated.

The authority's responsibility does not end with authorization. It must continue to monitor the use of AI during performance, ensuring that the technologies employed do not generate unlawful, discriminatory, or otherwise harmful effects.

6. Technological innovation, technical evaluation, and judicial review

The decision also highlights that technological innovation can be taken into account in the context of the most economically advantageous tender. However, the introduction of AI must still comply with the objective evaluation criteria set out in Article 108 of the Code and is subject to the limits of judicial review.

Case law – such as TAR Lombardia-Brescia Decision No. 1039/2016 – holds that technical evaluations by tender committees are generally not subject to judicial scrutiny unless they are manifestly illogical, unreasonable, or based on factual errors. AI must therefore be assessed with the same rigor applied to other components of the bid: it is neither inherently superior nor inherently suspect.

7. Conclusions

Council of State Decision No. 8092/2025 confirms that integrating AI into public procurement is not a unilateral decision by the economic operator. It requires informed authorization and continuous oversight by the contracting authority. The administration remains solely responsible for any automated processes used in the contract and must therefore be able to understand, evaluate, and monitor the proposed technologies.

The challenge ahead will be to build the internal expertise necessary to manage this technological shift, balancing innovation with the principles of transparency, data protection, security, and accountability that define public administration. AI can add value to public tenders, but only within a clear regulatory framework and under the supervision of an authority capable of truly governing how these systems operate.

AI and preventing insurance fraud from automatic identification to compliance

Author: *Giacomo Lusardi*

AI is perhaps the most advanced tool in the fight against insurance fraud. Machine learning and generative AI can analyse large volumes of heterogeneous data (claims, social media, open data) and identify anomalous patterns that could indicate attempted fraud.

But AI is also being used to trick consumers. One example is the "ghost broker" scam. Websites and chatbots, using generative AI, simulate real agencies and produce fake policies, managing to deceive even the most experienced users.

AI-based anti-fraud systems can cross-reference historical, behavioural and biometric data, detecting anomalies in real time, such as multiple requests from different individuals with identical or similar data, false documents generated by AI or serial claims. AI also helps automate cross-checking between public and private databases, drastically reducing identification times and the risk of human error.

The most advanced applications in the field of anti-fraud include:

- Computer vision techniques to compare images of claims and validate the authenticity of the damage detected.
- Behavioural biometrics to detect inconsistencies in user interaction patterns.
- NLP (Natural Language Processing) to analyse the content of claims and identify linguistic inconsistencies that could potentially indicate fraud.
- Automatic cross-checking of medical records, testimonies and satellite data to validate the veracity of catastrophic events.

The role of the AI Act and the institutional response

TThe EU has adopted Regulation (EU) 2024/1689, known as the AI Act, to address the challenges posed by the rapid evolution of AI. The AI Act represents the cornerstone of AI regulation and governance in the internal market. The primary objective is to protect safety and fundamental rights through a risk-based approach. But the AI Act also has significant implications for the insurance sector.

Some AI systems, like those for biometric recognition, credit scoring or emotion recognition, are classified as "high risk" and have to comply with strict requirements in terms of risk assessment, transparency, traceability and human oversight.

AI systems used to detect financial fraud don't automatically fall into this category, but they could if they're integrated with other features or systems already considered high risk. This paves the way for closer scrutiny of both the tools used by insurance companies and those misused by fraudsters.

The AI Act also introduces specific rules for systems that generate synthetic content like deepfakes or realistic imitations of people, which are increasingly common in ghost broker scams and identity theft. Suppliers and users of these systems will have to ensure transparency, affix digital watermarks and provide adequate documentation, increasing the traceability of any illegal uses. Member states and competent authorities now also have to monitor high-risk systems and coordinate at European level.

These developments are an opportunity for national authorities like IVASS to take a proactive role, not only in supervising the sector, but also in preventing and suppressing AI abuse in the insurance industry. The Insurance Supervisory Authority's reports pay close attention to attempts at fraud against policyholders.

Fraud attempts especially affect the motor vehicle liability insurance sector. It's an extremely large market, divided among many operators, many of whom operate online. But fraud attempts could affect all sectors of the insurance business, particularly those with a high incidence of serial fraud.

The Insurance Supervisory Authority hasn't yet expressed its opinion on the use of AI in systems designed to identify fraudulent situations. But the it is considering using AI for more routine activities, such as managing consumers' complaints.

Using AI skilfully and appropriately could significantly improve insurance products in general. It could encourage companies to offer services that users want. For example, assistance policies, where resources freed up in the serial and documentary management of contractual positions could be used for personal services.

This is a long process that requires an approach different from the world of services. But it seems industry regulations are already setting the path to follow by imposing a balanced cost/benefit ratio for insurance product customers.

Legal design tricks

Little tips to use legal design in your daily activities

Trick #13: human in the loop: Why the legal designer is essential for AI (and not the other way around)

Everyone talks about AI.

Few remember that without skilled humans, even the most advanced algorithm can make... poor decisions.

AI accelerates document production, but speed without oversight = regulatory chaos.

Legal documents aren't just texts: they are **experiences** that influence decisions, risks, costs, and trust.

This is where the **Legal Designer** comes in.

Author: *Deborah Paracchini*

What is the role of the legal designer?

The Legal Designer:

- **structures legal processes**
- **organizes and simplifies information**
- **translates complexity into actionable, understandable choices**

The goal? To produce documents that are clear, coherent, and usable.

This is exactly the spirit of **Human in the Loop:** the human remains in the AI decision cycle, intervening at different stages of design, development, and deployment to improve performance and reduce risks.

What does "Human in the Loop " mean?

Human in the loop means a person intervenes:

- in AI design
- in data preparation
- in output review
- in risk governance

Because only someone who understands the law can recognize when AI makes legal mistakes.

Let's explore the key contributions of the Legal Designer in enabling Human in the Loop.

Curating data for AI

AI "learns" from the data we provide.

If documents are ambiguous, redundant, inconsistent, or full of jargon... the output will be equally flawed!

The Legal Designer:

- simplifies and clarifies legal texts during pre-processing
- ensures legal meaning is clear
- removes ambiguity, errors, and unnecessary jargon
- creates clear and consistent information structures

The goal? To provide clean, AI-ready legal data.

Example: standardizing contract templates improves AI responses to commercial team requests.

Building a smart knowledge base

Without structure, even the best AI becomes a chaotic search engine.

The Legal Designer:

- organizes laws, policies, and clauses in a navigable way
- defines taxonomies and information pathways
- links information across legal, compliance, and business teams

The goal? To create a shared, coherent, and up-to-date knowledge base.

Example: a corporate repository of clauses with pre-defined selection criteria allows AI to suggest pre-approved texts, reducing risk and increasing speed.

Monitoring Risks and Output Quality

Technology moves fast. Law reflects.

The Legal Designer:

- sets limits and controls
- identifies ethical and legal risks
- reviews and corrects outputs (preventing hallucinations)

The goal? To avoid "legal hallucinations" and ensure compliance.

Example: a clear review workflow for AI-generated contracts prevents "fantasy legal outputs."

Putting people at the center

A tool nobody uses is wasted investment.

The Legal Designer:

- designs the legal user experience
- makes AI interactions clear and intuitive
- communicates when, how, and why to use the tool

The goal? To adopt sustainable AI that puts people first and technology at their service.

Example: simplified guides for employees using AI systems entail fewer requests to Legal, more autonomy for everyone.

In summary

AI **does not replace** human expertise. It amplifies the value of those who know how to guide it.

In this context, the Legal Designer is:

- curator of quality
- knowledge architect
- risk guardian
- translator between law, technology, and people

The Legal Designer is not a bystander in the technological transformation. They are a protagonist: bridging, guiding, and acting as the critical conscience of intelligent systems.

Without **Human in the Loop**, AI isn't intelligent: it's just automatic.

Did you know?

Internal studies and Big Tech experiments show removing 20% of unnecessary text from contracts increases AI accuracy by up to 30% in suggested clause modifications.

Want to design an AI-ready ecosystem ?

Start with **data**. Start with **people**.

And start with someone who can speak to both: the **Legal Designer**.

Legal tech bytes

Expert insights on the latest trends and innovations

Author: *Tommaso Ricci*

Generalist VS. Specific: Understanding which Legal tech tools to buy

The legal technology market is undergoing a fascinating evolution. As the sector matures and moves from pilot projects to production deployments, two distinct categories of tools are emerging, each serving different purposes and requiring different expertise levels. Understanding this spectrum is essential for legal teams making strategic technology investments.

The Swiss army knife: Versatile horizontal solutions

Consider the Swiss army knife: a remarkably useful tool that combines multiple functions in a single package. It can cut, open bottles, tighten screws, and perform dozens of other tasks competently. For most everyday situations, it is more than adequate. The same applies to horizontal LegalTech solutions currently dominating the market.

These platforms are designed to serve a broad range of legal tasks across practice areas. They excel at supporting drafting activities, providing document analysis capabilities, enabling side-by-side comparison during due diligence reviews, and offering quick access to source materials when verifying citations in legal opinions. Their strength lies in versatility: a single platform can assist with contract review in the morning and legal research in the afternoon.

The proliferation of these generalist solutions is partly driven by market dynamics. Investors backing LegalTech ventures expect scalability and returns on investment. A tool that addresses multiple use cases across different practice areas can reach a larger market faster than a highly specialized solution. This economic reality has shaped the current landscape, producing platforms that aim to be helpful across the entire spectrum of legal work.

According to industry data, approximately 70% of LegalTech investment remains concentrated in Anglo-centric markets, with most funding flowing toward platforms offering broad horizontal capabilities. These solutions have become the default starting point for many legal departments exploring AI adoption, precisely because they require minimal configuration and can demonstrate value across multiple workflows relatively quickly.

The scalpel: precision vertical solutions

No surgeon would perform an operation with a Swiss army knife. When precision matters and the stakes are high, professionals reach for the scalpel: a tool designed for one purpose, executed with extraordinary precision.

A parallel market is emerging in LegalTech: hyper-vertical solutions that do not attempt to solve a hundred problems, but instead solve one problem exceptionally well. These tools are built by specialists for specialists, often founded by lawyers who spent years practicing in specific domains and understood precisely what was missing from existing workflows.

Vertical LegalTech platforms are designed around the actual work patterns of practitioners in specific fields. A competition law AI platform understands merger control filing requirements across jurisdictions. A real estate technology solution integrates with land registry systems and knows property transaction workflows intimately. An intellectual property tool can navigate patent classification systems with the precision that generalist platforms simply cannot match.

The key differentiator is domain depth. Vertical solutions are trained on specialized datasets, integrate with industry-specific systems, and are designed around workflows that require genuine expertise to understand. They do not just process legal text: they understand the context, the regulatory requirements, and the professional standards that govern specific practice areas.

The expert hands factor

Here lies a crucial distinction: while horizontal tools are designed for broad accessibility, vertical solutions often require expert hands to unlock their full potential. A scalpel in untrained hands is merely a sharp blade. In the hands of a surgeon, it becomes an instrument of precision.

Vertical LegalTech tools assume domain knowledge. They speak the language of their specific practice area, use terminology that practitioners understand, and produce outputs calibrated to professional expectations. This is not a limitation but a feature: by assuming expertise, these tools can operate at a higher level of sophistication and deliver results that generalist platforms cannot replicate.

The trade-off is clear. Horizontal solutions lower the barrier to entry and can be deployed broadly with minimal training. Vertical solutions demand more from their users but reward that expertise with superior precision and workflow integration.

Strategic implications for legal teams

Neither category is inherently superior. The choice depends on workflow requirements, risk tolerance, and the nature of the work being performed.

For general productivity enhancement, communication management, and routine document analysis/editing, horizontal solutions offer immediate value with minimal friction. They are ideal for tasks where good enough is sufficient and where the cost of imprecision is low.

For specialized practice areas where accuracy is paramount, regulatory compliance is complex, and professional standards demand precision, vertical solutions provide capabilities that horizontal tools cannot match. The additional investment in learning specialized platforms pays dividends through higher accuracy, better workflow integration, and outputs that require less professional review and correction.

Many organizations are finding success with hybrid approaches: deploying horizontal platforms for general productivity while investing in vertical solutions for their core practice specialties. This portfolio approach acknowledges that different tools serve different purposes and that attempting to force a single solution to handle all use cases often results in suboptimal outcomes across the board.

Looking ahead

As LegalTech matures, both horizontal platforms and vertical solutions are becoming better. Legal professionals should choose technology based on their specific needs: general tools for broad routine tasks, specialized ones for precision work. Success depends on knowing when to use each tool.

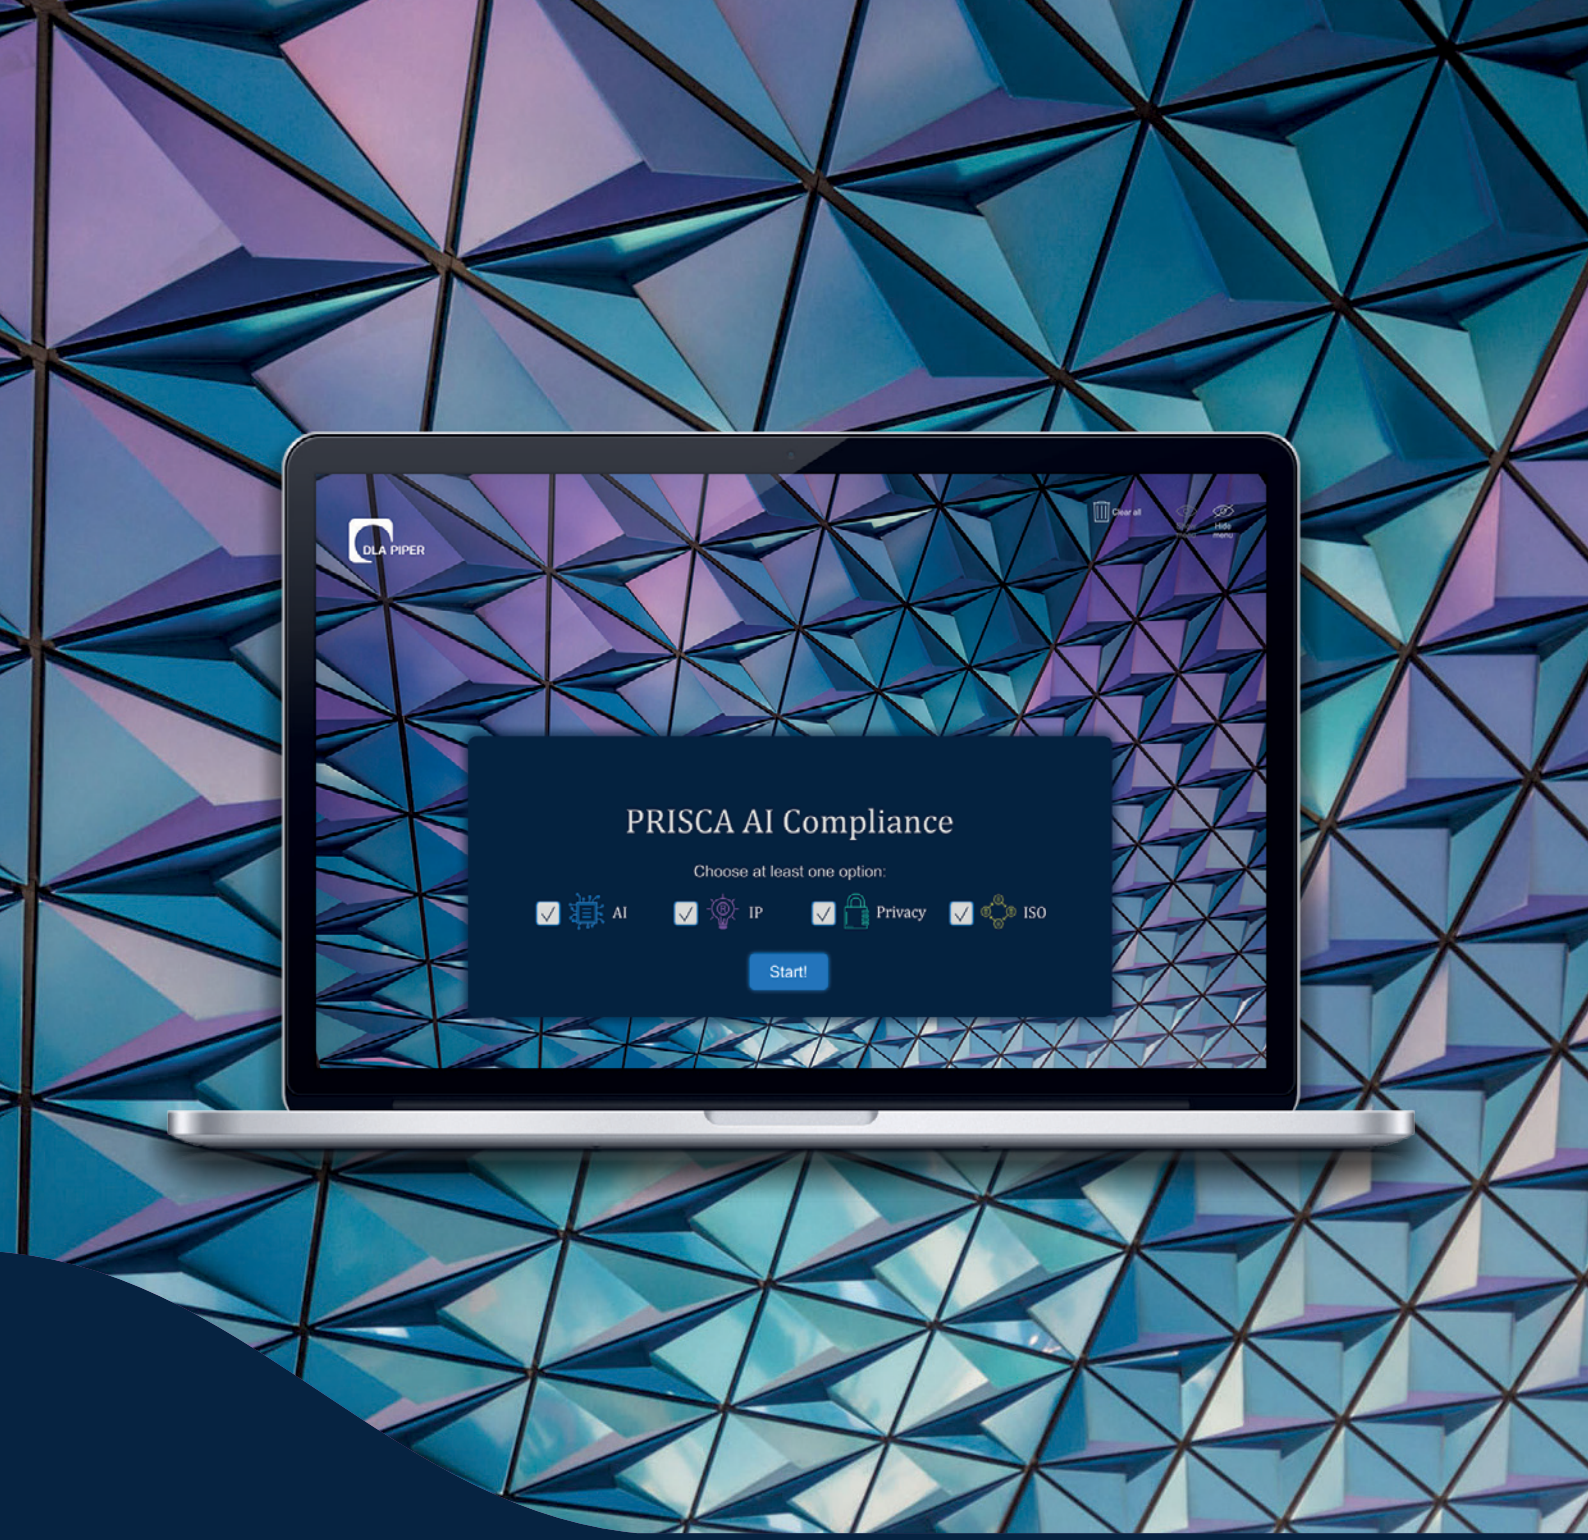

Prisca AI Compliance

Empowering Legal Compliance in the Age of Artificial Intelligence

Is your business ready to embrace the opportunities of AI, but worried about legal risks?

Introducing PRISCA AI Compliance by DLA Piper lawyers, a cutting-edge tool to assess your AI solutions' compliance with laws and ISO standards.

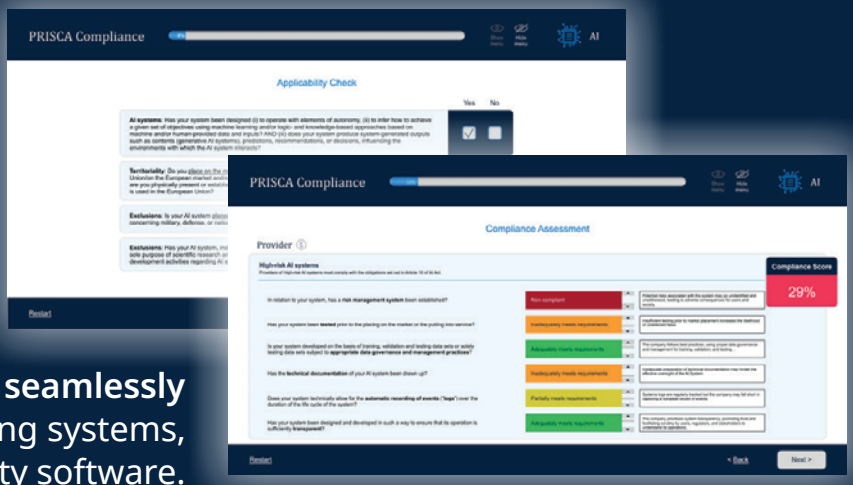

PRISCA AI Compliance **seamlessly integrates** into your existing systems, with no need for third-party software. It's available in **English** for global use.

Our unique weighted scoring algorithm generates a **compliance score** and an **easy-to-read report**. It highlights compliance with laws (privacy, IP, AI) and ISO standards.

Whether you're a user, provider, importer, or distributor of AI solutions, PRISCA AI Compliance supports your operations in **complying with regulations**.

Scan the QR Code to watch the video

Contact us for a demo:

giulio.coraggio@dlapiper.com
alessandro.ferrari@dlapiper.com
gualtiero.dragotti@dlapiper.com
elena.varese@dlapiper.com

Scan this qr code to access all issues of Diritto Intelligente

Contacts

Giulio Coraggio
Partner
Head of Intellectual Property and Technology, Italy
T +39 02 80 618 1
giulio.coraggio@dlapiper.com

Gualtiero Dragotti
Partner
Global Co-Chair, Patent Group
T +39 02 80 618 1
gualtiero.dragotti@dlapiper.com

Alessandro Ferrari
Partner
Head of Technology Sector, Italy
T +39 02 80 618 1
alessandro.ferrari@dlapiper.com

Roberto Valenti
Partner
Head of Life Sciences Sector, Italy
T +39 335 73 66 184
roberto.valenti@dlapiper.com

Elena Varese
Partner
Co-Head of Consumer Good, Food and Retail Sector, Italy
T +39 02 80 618 1
elena.varese@dlapiper.com

Ginevra Righini
Partner
T +39 02 80 61 863 4
ginevra.righini@dlapiper.com

Marco de Morpurgo
Partner
Global Co-Chair, Life Sciences
T +39 06 68 880 1
marco.demorpurgo@dlapiper.com

Alessandro Boso Caretta
Partner
T +39 06 68 880 1
alessandro.bosocaretta@dlapiper.com

dlapiper.com

Dec 18 2025 | A31340-2