Clifford Chance

BRIEFING

Making AI transparency work:

Four implementation questions for the Article 50 AI Code of Practice

22 May 2026

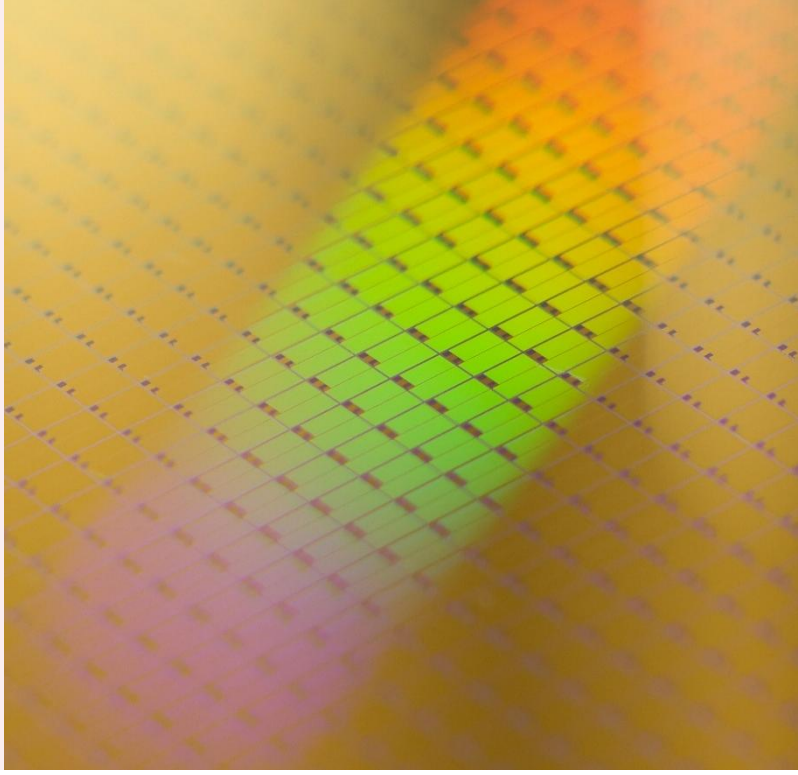

Introduction

The European Commission's ("**Commission**") voluntary Code of Practice on Transparency of AI-generated Content (the "**Code**") aims to facilitate the effective implementation of the transparency obligations under Article 50 of the EU Artificial Intelligence Act (the "**AI Act**").

Following the publication of an initial draft and stakeholder consultations, on 3 March 2026, the Commission released a second draft of the Code (the "**Draft Code**").

The Code is intended to provide practical guidance on how organisations can demonstrate compliance with requirements to apply machine-readable marks to AI-generated and AI-manipulated content, and to apply visible, human-readable labels to content that falls within the AI Act's definition of "deep fakes" and certain AI-generated text, ahead of those obligations generally becoming applicable on 2 August 2026 (subject to the EU's forthcoming Digital Omnibus on AI establishing a transitional period for certain Article 50(2) AI Act marking obligations).[1]

We have been tracking the development of the Code closely. Building on the substantial work undertaken by the Commission, the AI Office and the AI Board, we recently convened a roundtable discussion with industry, supported by Amazon, to debate and discuss the Draft Code. The discussion brought together a diverse range of stakeholders from across the AI value chain, spanning the tech, education and training, creative, industrial and entertainment sectors. Participants included European-headquartered companies, spanning large operators, scale-ups and SMEs, and providers of AI tools, digital content generation technologies and online services, including Black Forest Labs, Colossyan, Photoroom and Synthesia.

This article draws on the collaborative spirit of that discussion, highlighting four key implementation issues identified by participants. Our intention is

[1] The 4-month transitional period likely to be agreed under the Digital Omnibus on AI for labelling requirements (Article 50(2) AI Act) would only apply to general-purpose AI systems placed on the market *before* 2 August 2026. Importantly, any general-purpose AI system placed on the market on or after 2 August 2026 would be immediately subject to the Article 50(2) labelling requirements, with no transitional grace period available. In addition, unlike the GPAI Code of Practice, the AI Act does not expressly require the Article 50 Code of Practice to be adopted ahead of the date of application of the relevant obligations, and the process for its adoption is correspondingly less prescriptive.

to inform the ongoing debate and support the successful drafting of a Code that is workable, proportionate and fit for the future:

1. **Feasibility and proportionality of imperceptible watermarking requirements**, across content modalities and deployment contexts, given the limitations in the reliability and maturity of available watermarking and detection approaches, the integration of third-party and upstream models, and the compute, infrastructure and operational burdens associated with implementing and maintaining multi-layer marking approaches, particularly for scale-ups and SMEs.
2. **Prescriptiveness of testing, verification and monitoring obligations**, which go beyond the outcomes-based requirements set out in Article 50 AI Act and could give rise to ongoing compliance expectations that may be difficult to scale, particularly for SME providers.
3. **Scope of personnel training requirements within the Draft Code**, including how such measures can be applied proportionately and in a way that meaningfully supports transparency objectives.
4. **Breadth of the proposed labelling obligation for deep fakes and certain AI-generated or AI-manipulated content**, and the need for clearer guidance on when synthetic content meets the "deep fake" threshold. Without a clear distinction between benign or non-deceptive uses of synthetic content and content that is genuinely false or misleading, there is a risk that broad and undifferentiated labelling could lose explanatory value over time and, in turn, weaken rather than reinforce consumer trust.

Background

Article 50 AI Act sets out the transparency requirements applicable to providers and deployers of certain AI systems, focusing on ensuring that individuals are properly informed when they are interacting with, or exposed to, AI-generated or AI-manipulated content.

The Draft Code allocates responsibilities between (i) providers, to enable identification and detectability of AI outputs, and (ii) deployers, to facilitate labelling of deep fakes and AI-manipulated text.

In brief, Article 50 AI Act introduces baseline transparency requirements across four scenarios:

- **AI systems interacting with natural persons - Providers** must ensure individuals are informed when they are interacting with an AI system, unless this is obvious in context.[2]
- **Synthetic content generated or manipulated by AI systems - Providers** must ensure AI-generated or AI-manipulated audio, image, video or text outputs are marked in a machine-readable form and

[2] The Draft Code does not address Articles 50(1) or 50(3) AI Act. The Draft Code sits alongside the Commission's draft Article 50 Transparency Guidelines, published in May 2026. The Guidelines aim to serve as non-binding practical guidance to assist competent authorities, as well as providers and deployers of AI systems, in relation to Article 50 AI Act. They cover the scope of application, relevant legal definitions, the transparency requirements, and related exceptions.

detectable as such. Solutions must be effective, interoperable, robust and reliable, subject to technical feasibility and defined exclusions.

- **Emotion recognition and biometric categorisation systems - Deployers** must inform individuals who are exposed to the operation of such systems, subject to specified law-enforcement carve-outs.
- **Deep fakes and AI-generated text on matters of public interest - Deployers** must disclose certain deep fakes and certain AI-generated or AI-manipulated text, subject to defined exceptions.

The first draft of the Code, published on 17 December 2025, marked an initial attempt to translate the outcome-based transparency obligations in Articles 50(2), (4) and (5) of the AI Act from legal requirements towards operational application. Its reception, however, made clear that the approach taken in that draft did not yet provide a sufficiently workable or scalable framework for providers and deployers across diverse systems and use cases. In particular, several provider-level measures included in the first draft – such as requirements relating to provenance chains and functionality to support deployer disclosures for deep fakes – have since been refined, removed or made optional, reflecting a reassessment of how Article 50 AI Act's requirements can be implemented in practice.

The current Draft Code builds on this revised approach by setting out more detailed and structured proposals, including in relation to marking techniques, detection mechanisms and disclosure practices, and by outlining how considerations of technical feasibility and proportionality might be applied in practice.

At the same time, the second draft brings greater clarity to where further calibration may still be needed to ensure that Article 50 AI Act's outcomes-based obligations are implemented in a way that is both operationally workable and scalable across the AI value chain.

In particular, the measures contained in the Draft Code are significantly more prescriptive than other voluntary frameworks developed under EU legislation, such as those under the GDPR or the Digital Services Act. Those frameworks typically focus on setting high-level objectives and outcomes, giving organisations flexibility in how they design and implement compliance in practice. By contrast, the Draft Code takes a much more operational and technical approach, in parts resembling a standard-setting exercise. While the Code is intended as a flexible tool, there remains uncertainty as to how it will be treated in practice by national regulatory authorities, particularly for providers and deployers who choose not to sign. Against that background, this article focuses on four issues that warrant further consideration to improve feasibility, scalability and long-term effectiveness.

Issue 1 - Watermarking in practice (Sub-Measure 1.1.2 and Measure 2.3)

1. Requirements under the Draft Code

While Article 50(2) and (5) AI Act set outcomes based requirements for marking, detectability and clear disclosure of AI-generated and AI-manipulated audio, image, video and text content, with flexibility as to technical implementation, the Draft Code envisions a more prescriptive model. In particular, Sub-Measure 1.1.2 of the Draft Code requires imperceptible watermarking as a default element of compliance and

Measure 2.3 of the Draft Code sets out detailed expectations around verification and detection, including how results should be communicated to users, the techniques used, and the provenance of the detected content.

2. Industry perspectives

In discussions with industry, a number of practical concerns were raised regarding the feasibility and proportionality of the proposed watermarking requirements:

(a) **Technical feasibility across modalities**. Imperceptible watermarking raises feasibility questions across a range of content types and deployment contexts. Practical considerations were raised that current state-of-the-art approaches to watermarking struggle to ensure persistence and reliable detection for some modalities, especially at scale, and that this limitation may not be apparent to users of the forensic detectors which Measure 2.2 requires to be developed. As a result, it is not clear that mandating imperceptible watermarking and detection mechanisms can consistently support Article 50 AI Act's transparency objectives, or whether such requirements risk generating inaccurate or ambiguous signals through false positives that may be misleading to users and downstream actors.

For example, text and code outputs were identified by participants as presenting acute challenges. Unlike image and video content, text and software code are inherently malleable and frequently subject to editing, re-formatting, translation or partial reuse. Software code constitutes deterministic machine instructions, and concerns were raised that even minor changes from the watermarking process may compromise functionality, introduce errors and security vulnerabilities.[3] Some stakeholders also flagged the need to consider data protection and tracking risks associated with certain persistent marking techniques.

These concerns reflect the broader question – also implicit in the Commission's draft guidelines on Article 50 AI Act (the "**Draft Guidelines**") – of how Article 50 AI Act's technical criteria should be applied in the context of evolving and uneven technical maturity across modalities.

(b) **Two-layer marking**. Related questions were also raised about the feasibility of implementing two-layer marking requirements in practice, particularly where AI systems integrate multiple third-party or upstream models. It was discussed that in such environments, providers often operate within complex AI supply chains and have limited technical control over how marking is applied across different model components, making it difficult to apply layered watermarking solutions consistently across modalities, providers and deployment contexts, especially where responsibilities are distributed across the value chain.

(c) **Proportionality and cost**. Stakeholders also highlighted proportionality concerns relating to cost and operational impact, including potential competitive effects on the EU market. Implementing watermarking solutions, particularly for video, multi-modal or real-time generative content (e.g., chatbots), can

[3] The Draft Guidelines clarify that outputs generated in the form of a short sequence of numbers, symbols or letters, such as source code, does not fall into scope of Article 50(2) AI Act. This is not reflected in the Draft Code.

entail significant infrastructure and compute costs. For example, video watermarking has been shown to lead to a roughly 20% increase in bandwidth demand in some implementations, which directly translates into content delivery costs that must be absorbed by the company or consumer. These costs may have a more significant impact for SME providers, which raised questions about whether a prescriptive technical baseline could, in this case, create disproportionate operational burdens without delivering commensurate gains in transparency and consumer trust. There was notable interest in how regulators might implement a risk-based approach in interpreting and enforcing transparency requirements.

(d) **Global regulatory fragmentation**. For companies operating globally, there is also a concern that the Draft Code's mandatory marking requirements go further than existing approaches in other major jurisdictions, such as California, South Korea, China and India. Notably, participants explained that no other major jurisdiction currently requires either watermarking of free-form, non-containerised text or a prescriptive two-layer marking approach. This divergence raises questions about the administrative and engineering burden of developing jurisdiction-specific solutions that cannot be deployed consistently across markets.

Although national regulators are expected to assess compliance primarily by reference to the binding obligations set out in the AI Act, there was shared industry concern that the Code could, in practice, come to serve as an important reference point for interpreting what constitutes effective, interoperable, robust and reliable technical solutions – particularly in areas where common standards and accepted benchmarks are still emerging. In that context, ongoing dialogue with the Commission on how emerging standards and techniques develop over time will be critical to the effective and durable implementation of the Code's transparency objectives.

Issue 1: key considerations and views arising from the industry roundtable

The Code should emphasise outcomes based and standards driven transparency approaches rather than prescribed watermarking techniques.

3. Key considerations

The final Code should make clearer that imperceptible watermarking is only one voluntary technique that companies may choose to implement, rather than a mandatory default baseline across all outputs. In particular, Sub-Measure 1.1.2 should be reframed to reflect the outcomes-based approach of Article 50(2) AI Act, allowing signatories to demonstrate compliance through any combination of techniques that meet Article 50 AI Act's four criteria (i.e., effectiveness, interoperability, robustness and reliability), taking into account technical feasibility, cost, and state of the art – consistent with the Draft Code's own focus on these considerations.

In doing so, the Code should more explicitly recognise that standards-driven, metadata-based provenance approaches (such as C2PA), complemented by other equivalent mechanisms where appropriate, may be sufficient to meet the AI Act's criteria. While Measure 1.1 of the Draft Code leaves open the possibility of single marking approaches in the future, there should be clearer acknowledgement that each machine-readable marking approach has inherent strengths and limitations, and that the criteria in Article 50(2) AI Act must be interpreted and applied in light of current technical feasibility.

If a mandatory watermarking baseline is nonetheless retained in the Code, its scope should be substantially narrowed. In particular, the existing

exception for very short text outputs should be extended to include all free-form text content, especially real-time outputs. Such text can instead rely on transparency disclosures under Article 50(1) and existing measures under Article 50(4) AI Act relating to text content of public interest.

Issue 2 - Operational demands of mandatory testing (Measure 4.2)

1. Requirements under the Draft Code

Article 50(2) AI Act requires providers to ensure that AI-generated and AI-manipulated outputs are marked and detectable using technical solutions that are effective, interoperable, robust and reliable as far as technically feasible but stops short of prescribing how providers should test or reassess those solutions on an ongoing basis. The Draft Code sets out more detailed mandatory expectations, with Measure 4.2 requiring providers to regularly test and monitor marking and detection solutions in real-world conditions, using adaptive and use-case-specific threat modelling. It further requires providers to update benchmarks and other measurement and testing methodologies, periodically re-evaluate detection thresholds, and to document and address compliance shortcomings or adversarial attacks reported by third parties.

2. Industry perspectives

Industry stakeholders were in broad agreement that ongoing testing and monitoring plays an important role in supporting transparency and reliability objectives and should generally be encouraged as best practice. However, many stakeholders expressed concern that the level of detail and prescriptiveness introduced by the Draft Code goes beyond the scope of Article 50 AI Act, with scalability and proportionality implications for certain providers. In particular, the combination of mandated regular testing, continuous monitoring and evaluation, and documentation expectations for third-party reports may represent a disproportionate burden for SMEs and in B2B use cases where transparency risks are often reduced compared to public or consumer-facing contexts.

Issue 2: key considerations and views arising from the industry roundtable

The Code should frame testing, verification and monitoring measures as voluntary best practice guidance, rather than compliance requirements.

3. Key considerations

The Code should encourage testing, verification and monitoring measures that support the effective operation of marking and detection solutions, but not seek to establish mandatory standalone or ongoing compliance obligations that are not expressly addressed in Article 50 AI Act. There remains uncertainty as to how the detailed operational measures set out in the Draft Code may be treated in practice by regulators when assessing compliance – whether they will be regarded as non-binding indicators of good practice, or as expectations of what constitutes appropriate implementation of Article 50(2) AI Act.

Issue 3 - Resourcing AI transparency training (Measure 4.3)

1. Requirements under the Draft Code

Article 50(2) and (5) AI Act require providers to ensure that AI-generated or AI-manipulated outputs are marked and detectable, and that providers remain transparent about their capabilities, but neither provision imposes internal training obligations. However, Measure 4.3 of the Draft Code requires signatories to make proportionate efforts to provide appropriate

training to personnel responsible for ensuring that the marking and detection measures under the Draft Code are effectively implemented.

2. Industry perspectives

There was agreement among roundtable participants that the Draft Code's requirement for personnel training extends beyond the scope of Article 50 AI Act, although some industry participants were less concerned by the administrative burden posed by Measure 4.3 due to its emphasis on proportionality. Industry participants generally acknowledged the importance of training and internal awareness of transparency measures as part of responsible AI governance and adoption, and many organisations already invest in internal education and cross functional coordination to ensure that technical, legal and operational teams understand how AI legal obligations apply in practice. Some participants mentioned that Article 4 AI Act already addresses AI literacy and is being operationalised by providers in scope of the AI Act – a separate mandatory training requirement for transparency was therefore considered beyond the scope of Article 50 AI Act.

3. Key considerations

To avoid regulatory uncertainty and reduce the administrative burden on organisations, the Code should make clear that the provisions on training and awareness in relation to transparency is intended to support the wider, existing obligation for providers of AI systems to ensure a sufficient level of AI literacy, rather than forming a standalone compliance obligation under Article 50 AI Act.

Issue 3: key considerations and views arising from the industry roundtable

The Code should recognise and encourage personnel training as a part of wider AI literacy efforts, rather than framing it as a standalone compliance obligation.

Issue 4 - Drawing the line on deep fakes

In parallel with our broader engagement on the provider focused transparency obligations under Article 50(2) AI Act, we have also engaged with industry on the Draft Code's separate provisions imposing disclosure obligations on deployers in relation to deep fakes and certain AI generated or AI-manipulated text.

1. Requirements under the Draft Code

Article 50(4) AI Act requires deployers to disclose when image, audio or video content has been generated or manipulated by AI and constitutes a deep fake, and similarly where AI-generated or AI-manipulated text is published in order to inform the public on matters of public interest. Crucially, the AI Act defines a "deep fake" by reference to content that would "falsely appear to a person to be authentic or truthful". The Draft Code seeks to operationalise this requirement through a more detailed labelling framework but, in so doing, raises a more fundamental question about where the line should be drawn in practice between innocuous synthetic content and content that is genuinely deceptive.

2. Industry perspectives

Industry participants cautioned that a broad interpretation of "deep fakes" would risk capturing a wide range of benign and consent-based uses of synthetic media – such as synthetic product imagery, AI-generated fashion models, authorised digital replicas, or AI-generated avatars used for internal company training – and subjecting them to the same prescriptive labelling requirements under the Draft Code as non consensual, deceptive or malicious deep fakes, notwithstanding the very different risks involved.

In such cases, the artificial nature of the content is typically apparent or expected, realism is an intended aesthetic rather than a tool of deception, and onward public dissemination may be limited or unlikely. Participants therefore viewed such cases, whether in B2B contexts or B2C contexts at a user's request, as presenting a reduced level of transparency risk compared to content designed to mislead, deceive or influence consumers or the public.

Against this backdrop, participants also raised concerns about usability and consumer trust – where "deep fake" labels are applied widely to synthetic content and without discrimination, there is a risk of "label fatigue", where disclosures become ubiquitous and lose their signalling value, making it harder for users to distinguish genuinely misleading or harmful content.

Participants also highlighted the practical challenges in allocating responsibility across the AI value chain, particularly where synthetic content is generated, modified and deployed to end-users by different operators.

These issues are particularly acute in light of the Draft Guidelines, which have prompted renewed debate as to how broadly the concept of "deep fake" may be interpreted in practice and how far labelling obligations could extend into routine or consent-based synthetic content.

Issue 4: key considerations and views arising from the industry roundtable

The Code and final Commission guidelines should clarify the scope of deep fakes by reference to a context-based threshold, such as harm or deception.

3. Key considerations

The Code, together with the Draft Guidelines, should focus on explaining how the concept of a "deep fake" should be applied in practice, including by clarifying that labelling obligations are triggered by risks of deception or harm, rather than by the mere fact that content is synthetically generated.

Such details could usefully provide illustrative examples of in-scope and out-of-scope use cases, helping distinguish genuinely misleading or harmful deep fakes from routine, benign or consent-based uses of AI-generated media.

In parallel, it is likely that national regulators will further refine their operational interpretation of "deep fakes" over time. There is an opportunity for these approaches to adopt a risk-based and context-sensitive assessment, taking into account factors such as the nature of the use case, the audience and dissemination context, and the distinction between commercial or professional deployments and content used in public interest or political settings.

Conclusion

Article 50 AI Act establishes a set of transparency obligations designed to ensure that individuals are appropriately informed when they are interacting with, or exposed to, AI-generated or AI-manipulated content. The Draft Code represents a further step towards the operationalisation of the AI Act's transparency obligations into practical guidance for both providers and deployers of AI systems.

However, discussions with industry stakeholders highlight the challenges inherent in translating outcomes-based legal obligations into detailed operational expectations that are workable across a wide range of systems, modalities and deployment contexts. Many of the questions raised are not points of principle, but of technical feasibility, proportionality and clarity,

particularly where prescriptive measures risk outpacing current technical capabilities or obscuring the deception-focused intent of the AI Act's deep fake provisions.

A further complexity is the ambiguous position the Code occupies within the regulatory framework – while formally voluntary, it may in practice function as a benchmark in supervisory assessments, including for non-signatories. This reinforces the importance of ensuring that the Code remains firmly anchored in the AI Act's scope and intent. The publication of the Draft Guidelines is an important step in shaping the emerging framework, but it does not resolve these underlying tensions – and, in some respects, brings them into sharper focus, particularly in relation to scope, proportionality and technical feasibility.

Against this backdrop, as the Draft Code and Draft Guidelines continue to be developed, a focus on outcomes, proportionality and practical application will be important to ensure that transparency measures remain meaningful for users while being workable for all providers and deployers in the European AI value chain. Further clarification through the evolving Code, the Draft Guidelines and national supervisory guidance – supported by concrete examples and contextual interpretation – has the potential to help bridge the gap between statutory objectives and operational reality.

Ongoing engagement with industry, regulators and other stakeholders will be key to ensuring that the transparency framework under Article 50 AI Act supports effective and proportionate implementation, strengthens user trust, and accommodates innovation across the EU's AI ecosystem.

Jonathan Kewley
Partner and Chair of the Global Tech Group, London

jonathan.kewley
@cliffordchance.com
+44 2070063629

Phillip Souta
Global Director of Tech Policy, London

phillip.souta
@cliffordchance.com
+44 2070061097

Monika Zdzieborska
Counsel, London

monika.zdzieborska
@cliffordchance.com
+44 2070061468

Eliot Cohen
Lawyer, London

eliot.cohen
@cliffordchance.com
+44 2070062966

Ryan Lin
Lawyer, London

ryan.lin
@cliffordchance.com
+44 2070061949

Anna Carrier
Head of EU Tech Policy, Brussels

anna.carrier
@cliffordchance.com
+32 25335048

Rita Flakoll
Global Head of Tech Group Knowledge, London

rita.flakoll
@cliffordchance.com
+44 2070061826

Alexander Kennedy
Knowledge Director – CE Tech Group, Paris

alexander.kennedy
@cliffordchance.com
+44 2070061949

This publication does not necessarily deal with every important topic or cover every aspect of the topics with which it deals. It is not designed to provide legal or other advice.

cliffordchance.com

Clifford Chance, 10 Upper Bank Street, London, E14 5JJ

© Clifford Chance 2026

Clifford Chance LLP is a limited liability partnership registered in England and Wales under no. OC323571. The firm's registered office and principal place of business is at 10 Upper Bank Street, London E14 5JJ. The firm uses the word "partner" to refer to a member of Clifford Chance LLP or an employee or consultant with equivalent standing and qualifications.

If you do not wish to receive further information from Clifford Chance about events or legal developments which we believe may be of interest to you, please either send an email to nomorecontact@cliffordchance.com or by post at Clifford Chance LLP, 10 Upper Bank Street, Canary Wharf, London E14 5JJ

Abu Dhabi • Amsterdam • Barcelona • Beijing • Brussels • Bucharest** • Casablanca • Delhi • Dubai • Düsseldorf • Frankfurt • Hong Kong • Houston • Istanbul • London • Luxembourg • Madrid • Milan • Munich • Newcastle • New York • Paris • Perth • Prague** • Riyadh* • Rome • São Paulo • Shanghai • Singapore • Sydney • Tokyo • Warsaw • Washington, D.C.

*AS&H Clifford Chance, a joint venture entered into by Clifford Chance LLP.

**Clifford Chance has entered into association agreements with Clifford Chance Prague Association SRO in Prague and Clifford Chance Badea SPRL in Bucharest.

Clifford Chance has a best friends relationship with Redcliffe Partners in Ukraine.