CLIFFORD

CHANCE

THE EU'S ARTIFICIAL INTELLIGENCE ACT:
WHAT DO WE KNOW ABOUT THE CRITICAL POLITICAL DEAL?

CONTENTS

THE EU’S ARTIFICIAL INTELLIGENCE ACT WHAT DO WE KNOW ABOUT THE CRITICAL POLITICAL DEAL?

The EU institutions have finally reached a political agreement on the EU's landmark Artificial Intelligence Act (AI Act), following the conclusion of their fifth and final round of negotiations late on 8 December.

This is a crucial milestone and a decisive step for the regulation of AI worldwide. The significance of this achievement should not be underestimated, no matter the work and efforts remaining to reach adoption of the AI Act.

As the key stakeholders work to finalise a complete, consolidated final text for approval, we share some initial thoughts and takeaways on some key aspects of the political agreement.

A word of caution: there is currently no complete, reliable final text. Our summary below is informed by the discussions and developments over the past months, as well as information from official publications and comments and statements made by those involved in the legislative process. Organisations will need to see the final text for certainty as to what the requirements and restrictions will be.

Executive overview of the political deal achieved on 8 December 2023

The political deal reached on 8 December is a critical milestone. Reportedly, the EU institutions have managed to come to an agreement on key and complex issues:

- Rules to regulate general-purpose AI models, involving a tiered system with rules for all general-purpose AI models, and additional rules for general-purpose AI models with systemic risks.
- A ban on real-time remote biometric identification in public spaces for law enforcement purposes, and exceptions to that ban.
- The other practices that are deemed particularly harmful and will therefore also be prohibited under certain circumstances, which includes expanding the list initially envisaged.
- The governance, supervision and enforcement of the AI Act, with different bodies and forums to be set up. These include an AI Office within the European Commission, which will play a key role.
- Numerous other important topics, such as:
  - the definition of an AI system
  - the scope of the AI Act, including exclusions, clarifying questions re Member State competencies, and interplay with sectoral legislation
  - the list and classification of 'stand-alone' high-risk AI systems, including a 'filter system', and the requirements for high-risk AI systems
  - the introduction of a requirement to carry out a prior fundamental rights impact assessment
  - measures to support innovation
  - specific provisions around free and open-source AI
  - the penalties and fines for infringing the AI Act, etc.

The reported acceleration of the planned entry into application of certain requirements of the AI Act, e.g., as regards prohibited practices and general-purpose AI, will also be business critical.

This is not the end of the road, and work remains to finalise and approve a complete and consolidated text.

Refresher: What is the EU's AI Act and what does it aim to do[1]?

- The EU's AI Act is the EU's proposal to regulate the placing on the market, putting into service and use of AI systems in and affecting the EU.
- It will set legally binding rules and requirements for AI systems and models, affecting actors across the entire AI value chain.
- It will have global reach, and will not be limited to EU-based businesses. For instance, third country providers placing on the market or putting into service AI systems in the EU will be captured. Beyond, it is expected to set the standard or at the very least a benchmark globally.
- It will provide strong enforcement measures. These will include significant, gradual fines: following the political agreement, they could reach €35 million or 7% of the global annual turnover for the most important infringements.
- It generally follows a risk-based approach, mainly articulated around: (a) the prohibition of unacceptable AI practices; (b) a broad set of rules and requirements for high-risk AI systems, that are at the core of the proposal; (c) specific transparency requirements to address specific types of AI systems or AI uses, e.g., those that interact with people such as chatbots, or deepfakes.
- Further, developments these past months have put a strain on this risk-based approach, in particular certain developments around the regulation of 'foundation models', general purpose AI and generative AI, as well as general principles for AI systems. We look at some of these aspects in more detail in this note.

Where do we stand on the question of foundation models and general-purpose AI?

One of the most debated topics this year, and which is in the spotlight worldwide[2], is the regulation of foundation models (the powerful models that notably support, and are the foundation for, other AI systems, including generative AI systems for instance), general-purpose AI, and generative AI.

In the weeks leading up to the fifth trilogue, we saw resistance to specific rules and requirements for foundation models from some EU countries – these argued in favour of alternative approaches such as mandatory self-regulation through codes of conduct, as well as European standards.

Following the start of the trilogues, numerous proposals and compromises were put forward and discussed on the regulation of foundation models and general-purpose AI, building on what the Council, and then to a greater extent the Parliament, had proposed as their respective negotiating positions for the trilogues.

The political agreement reached by the EU institutions is for the AI Act to establish a tiered system, a concept that was prominent in proposals and discussions in recent weeks. Based on available information, this system will focus on the notion of general-purpose AI, and in particular general-purpose AI models (rather than the specific notion

[1] For some further reading: (a) for our views and analysis on the European Commission's proposal of April 2021: The future of AI regulation in Europe and its global impact; (b) for our key takeaways following the European Parliament's negotiating position of June 2023: EU AI Act: Final negotiations can begin after European Parliament vote.

[2] See, for instance, Biden's Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence (for initial thoughts on what businesses need to know about the Executive Order, see here: What businesses need to know (for now) about the Biden Executive Order on AI), as well as the G7 'Hiroshima Process' International Guiding Principles and International Code of Conduct for Organizations Developing Advanced AI Systems.

of 'foundation models' itself). Under this system, all 'general-purpose AI models' would be subject to a set of requirements, and 'general-purpose AI models with systemic risks' would, in addition, be subject to further rules.

It will be key to see the detailed provisions in the definitive text when settled. That said, based on the information available to date, the requirements under the tiered system would appear to split somewhere along the following lines[3]:

(a) **Requirements for all 'general-purpose AI models'**

> The AI Act would include requirements such as: drawing up relevant technical documentation; drawing up and making available information and documentation to providers of AI systems who intend to integrate the general-purpose AI model in their AI system; putting in place a policy to respect EU copyright law; and making publicly available a summary on the content used for the model's training.
>
> There would likely also be transparency requirements regarding artificially generated or manipulated content, although probably as part of more general requirements.

(b) **Additional requirements for 'general-purpose AI models with systemic risks'**

> The AI Act would include rules on: model evaluation; the assessment and mitigation of possible systemic risks; the management of serious incidents and related corrective measures; adversarial testing; adequate cybersecurity protection; and possibly also requirements around reporting on energy consumption, as part of a push for enhanced emphasis on the environmental sustainability of AI.

Key features of this tiered system are expected to also include:

- The definition, and the criteria and procedure for the classification and designation, of general-purpose AI models with systemic risks.

  In this respect, one of the key reference points to determine whether a general-purpose AI model should be classified as one with systemic risks would be the amount of compute used for its training measured in floating point operations (or FLOPs). The European Commission's Q&A on the Artificial Intelligence Act, updated on 12 and 14 December 2023 ("**Commission's Q&A**"), seems to confirm that general-purpose AI models trained using total computing power greater than 10^25 FLOPs would in principle be considered as general-purpose AI models with systemic risks[4]. The Q&A provides its take on the appropriateness of that threshold, and it indicates that there would indeed be flexibility built into the system, including to enable the AI Office to update the threshold in the light of technological developments.

  It is also interesting to compare this with the thresholds in terms of computing power envisaged in Biden's Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, including as potential triggers for certain contemplated reporting requirements for dual-use foundation models or to

[3] It will also be interesting to see what developments there may be specifically on general-purpose AI systems.

[4] That threshold reportedly continues to be questioned if not challenged by some Member States following the political deal.

assess when a large AI model would be considered as having potential capabilities that could be used in malicious cyber-enabled activity (subject to further determination of those thresholds).

- Further process and transparency requirements, including possibly notification requirements for providers whose models meet the relevant requirements, and the publication of a list of general-purpose AI models with systemic risks.
- A central role for the European Commission regarding general-purpose AI models with systemic risks, including as regards their designation, and supervision and enforcement, and more generally mechanisms for the supervision and enforcement of general-purpose AI models.
- A framework for the development of codes of practice and harmonised European standards.
- Potential specificities for AI systems and AI models made available under free and open-source licences.

In addition, the final text of the AI Act may seek to further address the interplay between the requirements under the AI Act and those under the Digital Services Act, including as regards systemic risks for very large online platforms and very large online search engines, and watermarking.

What about remote biometric identification and facial recognition?

Another key debate has been around the scope of the ban the AI Act would impose on remote biometric identification.

As regards more specifically real-time remote biometric identification systems[5], the AI Act would prevent their use in publicly accessible spaces for law enforcement purposes, save in specific circumstances and subject to certain safeguards and conditions (that are reported as having been tightened compared with those in the Commission's proposal of April 2021 and the Council's General Approach of December 2022). According to some sources, the use of real-time remote biometric identification systems in publicly accessible spaces for law enforcement would be possible in connection with specifically defined law enforcement activities, including as related to a specified list of crimes[6], and subject notably to authorisation from a judicial or independent administrative authority (generally prior authorisation, but possibly with some exceptions).

This discussion also ties into a more general issue around the use of AI for law enforcement purposes. Other specificities and exceptions to the rules envisaged in the AI Act have reportedly been agreed.

Refresher: Prohibited practices – Parliament's negotiating position (June 2023)

In addition to real-time remote biometric identification in publicly available spaces, examples of prohibited practices added in the Parliament's negotiating position in June included the placing on the market, putting into service or use of:

- Biometric categorisation systems to categorise people according to sensitive or protected attributes or characteristics.
- AI systems to assess the risk of a person offending or reoffending or to predict the risk of occurrence or reoccurrence of a criminal or administrative offence based on profiling or assessing personality traits and characteristics.
- AI systems creating or expanding facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage.
- AI systems to infer emotions in the areas of law enforcement, border management, in the workplace and education institutions.

The Parliament also wanted to include the putting into service or use of AI systems for the analysis of recorded footage of publicly accessible spaces through post remote biometric identification systems, unless there is a pre-judicial authorisation and this is strictly necessary for the targeted search in relation to specific serious criminal offences.

[5] The documentation currently available appears somewhat confusing on the outcome concerning post-remote biometric identification. We do not address that issue here. The situation will need to be assessed in the light of the final text.

[6] The Commission's Q&A lists 16 crimes, and it also refers to other specified activities such as the targeted search for certain victims / persons and the prevention of threat to the life or safety of persons or response to the threat of a terror attack.

What is the status on the list of prohibited practices, and a possible trade-off between prohibited practices and high-risk AI systems?

The European Parliament, in its negotiating position in June 2023, had proposed to significantly expand the list of prohibited practices compared with the Commission's initial proposal.

Despite numerous points of contention, the EU institutions have managed to reach agreement on the list.

It seems that most of the practices suggested by the Parliament have made their way into the list, in some form or other. The list of prohibited practices would therefore now cover, in addition to certain situations of social scoring, manipulative practices, practices exploiting vulnerabilities, as well as certain forms of remote biometric identification systems[7]:

(a) Certain instances of predictive policing.

(b) Some biometric categorisation systems, using sensitive characteristics.

(c) Emotion recognition systems for use in the workplace or educational institutions.

(d) AI systems creating or expanding facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage.

In limited instances, the EU institutions appear to have agreed to add something that closely reflects the Parliament's June proposal. In most cases, however, the ban as envisaged by the Parliament has been made subject to caveats or additional conditions, limited in scope or moved, at least in part, to the high-risk category under Annex III.

Who's in charge at the EU level?

The question of the governance, supervision and enforcement of the AI Act has also been an important discussion point, with developments over time around different bodies (in addition to national competent authorities, which will notably be involved in the enforcement at a national level) such as: an AI Board, comprised of representatives of the Member States; an AI Office, within the European Commission; an Advisory Forum, with a balanced selection of stakeholders; and a Scientific Panel of Independent Experts.

We anticipate that all the bodies mentioned above would in principle be set up under the AI Act, with different roles and responsibilities. The AI Office would have a central role. It would notably be responsible for ensuring EU-level co-ordination, and for overseeing, and enforcing the rules on, general-purpose AI models.

Some expect the AI Office to become an international reference, being the first body globally to enforce binding rules on AI, and thereby further contributing to the global impact and reach of the EU's regime for regulating AI.

[7] As indicated above, the situation with respect to post-remote biometric identification seems unclear in the light of the information currently available.

What other topics were discussed in the political trilogues?

We have highlighted the areas above for this note but there are many other key topics that had to be discussed and agreed during the political trilogues to reach this political agreement. They include:

(a) **The very definition of AI system**, including in the light of the revised OECD definition published in November 2023 with which the proposed AI Act reportedly aligns. This is a key issue, going to the core of the scope and reach of the AI Act.

(b) **The scope of the AI Act**, including the exclusion of activities falling outside the remit of EU law, and notably the exclusion of systems developed or used (exclusively) not only for military purposes but also for defence purposes. There would also be a clarification concerning national security and Member States' competences in this area, one of the touchy discussion points during the trilogues.

(c) **The list and classification of systems as stand-alone high-risk AI systems under Annex III,** including the proposed 'filter system', and, of course, **the requirements applicable to high-risk AI systems**.

Stand-alone high-risk AI systems under Annex III would include certain AI systems in the fields of: biometric identification, categorisation and emotion recognition (to the extent not prohibited); certain critical infrastructure; education and vocational training; employment and workers management; access to and enjoyment of essential public and private services; law enforcement, migration, asylum and border control management; and administration of justice and democratic processes.

On the filter system, briefly: if one or more specific criteria are met – to take one example, the relevant AI system is intended to perform only a narrow procedural task of low complexity – there would be circumstances where an AI system coming within Annex III could nonetheless be considered not high-risk in that case, and hence potentially exempted from the related rules. A specific process would apply, in principle involving self-assessment by the provider, control procedures and possibly fines for the provider if it has sought to circumvent the requirements of the AI Act, based on reports during the trilogues. Also, there would be a general safeguard where the relevant AI system performs profiling of natural persons: such systems would always be considered to pose a significant risk of harm to the health, safety or fundamental rights, and hence be deemed high-risk.

(d) **The introduction of a requirement to carry out a prior fundamental rights impact assessment.** According to the Commission's Q&A, deployers that are bodies governed by public law or private operators providing public services, as well as operators providing high-risk AI systems, will be subject to the requirement to conduct a fundamental rights impact assessment. The European Parliament's press release specifically emphasises that this would apply also to the insurance and banking sectors.

(e) **A potential list of 'general principles' for AI systems.**

(f) **Further questions around the interplay with existing sectoral legislation.**

(g) **Greater transparency on the fact that content is artificially generated or manipulated.**

(h) **Measures to support innovation**, including re AI regulatory sandboxes as well as questions around the testing of high-risk AI systems in real world conditions outside of AI regulatory sandboxes.

(i) **Developments to address free and open-source AI.**

(j) **Specificities for SMEs, including start-ups.**

(k) **The enforcement measures and significant fines under the AI Act.**

What's next for the AI Act?

The EU institutions have come to a political agreement, but this is not the end of the road:

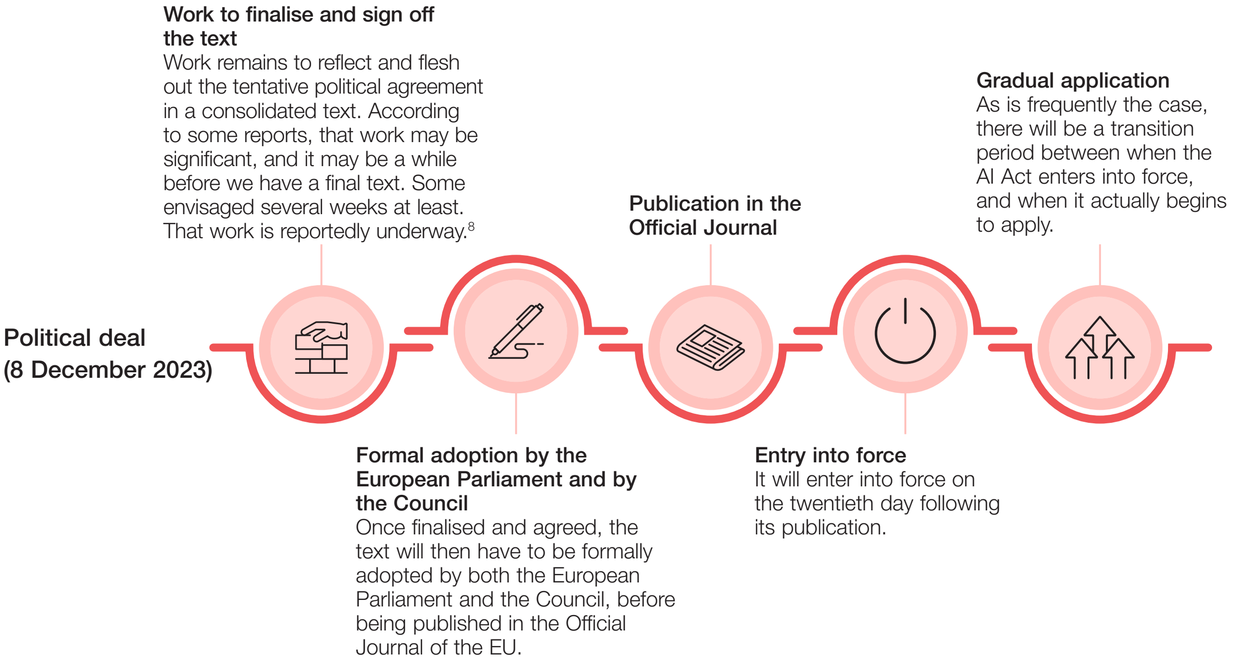

Political deal (8 December 2023)

Work to finalise and sign off the text
Work remains to reflect and flesh out the tentative political agreement in a consolidated text. According to some reports, that work may be significant, and it may be a while before we have a final text. Some envisaged several weeks at least. That work is reportedly underway.[8]

Formal adoption by the European Parliament and by the Council
Once finalised and agreed, the text will then have to be formally adopted by both the European Parliament and the Council, before being published in the Official Journal of the EU.

Publication in the Official Journal

Entry into force
It will enter into force on the twentieth day following its publication.

Gradual application
As is frequently the case, there will be a transition period between when the AI Act enters into force, and when it actually begins to apply.

**Important secondary legislation and other documents will also need to be adopted / prepared under the AI Act, including delegated acts and implementing acts.*

Bearing in mind, as regards the timing for the finalisation and adoption of the text, that the European elections of June 2024 act as a natural deadline.

When will the AI Act really start applying?

According to the EU institutions' press releases following the political agreement, and whilst different options had been contemplated, it seems that the transition period would generally be of **24 months following entry into force**, subject to exceptions for specific provisions. On this last point, and according to the European Commission's press release and the Commission's Q&A, it would seem that[9]:

- The rules on **prohibited practices** would apply after **6 months**;
- The rules on **general-purpose AI** would apply after **12 months**; and

[8] Reportedly, various technical meetings were already scheduled including several that took place the week of 11 December 2023 for instance, and the issue was also discussed in the context of a meeting on 15 December 2023 where the Spanish Presidency of the Council debriefed Coreper on the outcome of the trilogue. Discussions and developments will need to be closely monitored in the coming weeks.

[9] The final text is likely to include further nuances and specificities. Also, certain other provisions (e.g., re notifying authorities and bodies to be set up by the Member States, the EU-level governance structure to be implemented, the rules to be laid down by Member States with respect to the penalties applicable to infringements under their powers, certain documents to be prepared / adopted under the AI Act) may also be subject to specific timeframes and generally apply earlier to enable the implementation of the AI Act.

- **Obligations for high-risk AI systems defined in Annex II** (we understand this to mean AI systems classified as high-risk due to their being products or safety components of products coming under specific sectoral legislation listed in Annex II and being subject to a third-party conformity assessment under that legislation) would apply after **36 months**.

There will also be specific provisions to address AI systems that are already on the market or in use when the AI Act begins to apply.

If the above is confirmed in the final text, and assuming the AI Act is finalised and published in the next few weeks or months, this would notably mean that:

Some requirements could begin to apply in the second half of **2024**.	Other requirements could begin to apply in the first half of **2025**.	With the bulk of requirements, including requirements for high-risk AI systems[10], then kicking in during **2026**.

And in the meantime, what of voluntary initiatives such as the AI Pact?

There has also been a push these past months to develop voluntary initiatives and encourage businesses to anticipate certain requirements before they become law.

There have been important developments on the international stage, including the G7 'Hiroshima Process' International Guiding Principles and International Code of Conduct for Organizations Developing Advanced AI Systems.

In the EU more specifically, there is also the AI Pact. The AI Pact aims to foster early implementation by businesses of the AI Act, encouraging the sharing of processes and best practices as well as voluntary commitments to "anticipate and bridge the gap" before the AI Act actually begins to apply.

The European Commission launched a call for interest in November, with the ambition of starting to bring interested parties together in the first half of 2024. The AI Pact itself would be officially launched following the AI Act's formal adoption. The AI Pact is open to both EU and non-EU businesses, given the scope of the AI Act.

According to the statement issued by European Commission President von der Leyen following the political agreement, around 100 companies had already expressed their interest in joining.

A wider regulatory framework in the EU for AI

The AI Act is one part of the regulatory framework developing in the EU to address the specificities of AI.

There are a number of other associated EU legislative texts that have recently been adopted or that are under discussion, including the 2023 Machinery Regulation, the proposed AI Liability Directive, and the proposed revision of the 1985 Product Liability Directive. On 14 December 2023, a provisional political agreement was reportedly reached on the latter as well.

Beyond, many texts that are not AI-specific contain requirements impacting AI, including the General Data Protection Regulation, the Digital Services Act and cyber-related legislation, amongst others.

[10] Also potentially subject to some specificities, as flagged above.

What of the broader context?

The regulation of AI is a global issue. These past few months have seen increased international attention. Examples include:

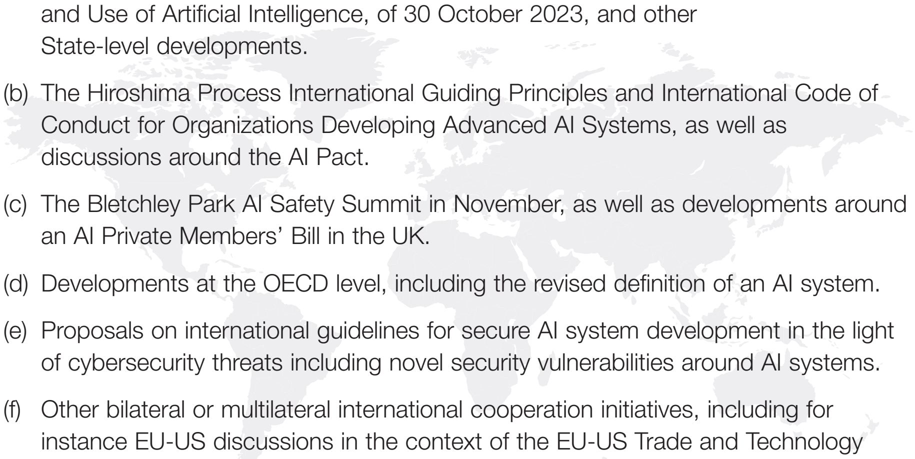

(a) Biden's Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, of 30 October 2023, and other State-level developments.

(b) The Hiroshima Process International Guiding Principles and International Code of Conduct for Organizations Developing Advanced AI Systems, as well as discussions around the AI Pact.

(c) The Bletchley Park AI Safety Summit in November, as well as developments around an AI Private Members' Bill in the UK.

(d) Developments at the OECD level, including the revised definition of an AI system.

(e) Proposals on international guidelines for secure AI system development in the light of cybersecurity threats including novel security vulnerabilities around AI systems.

(f) Other bilateral or multilateral international cooperation initiatives, including for instance EU-US discussions in the context of the EU-US Trade and Technology Council, announced cooperation between the EU and Latin America and Caribbean (LAC) countries, the EU-Canada Digital Partnership, developments in the context of the Global Partnership on Artificial Intelligence, and other such initiatives.

And whilst the regulation of AI can get caught up in important economic, (digital) sovereignty and geopolitical considerations, there is a need for a minimum level of international cooperation and alignment on some fundamentals.

What can organisations do now?

For businesses, it will be important to have the definitive version of the text to be able to fully assess the impact of the AI Act on their operations and strategy and to step-up preparation efforts whilst having the necessary visibility and legal certainty.

Nevertheless, there are steps that organisations can already be taking. Here are five examples:

1

Map AI plans and activities and assess the implications

It is important for organisations to understand and map how they are developing, providing and/or using AI within their businesses, and what their plans are in this respect.

From a regulatory standpoint, this enables organisations to more quickly identify the implications, under existing and upcoming rules, carry out the necessary gap analyses, and define and implement the necessary actions.

And looking specifically at the AI Act: whilst it is not yet law, and the final text is needed for the deep dive, important work can already be done – all the more so to the extent certain requirements may be kicking in as early as 2024.

2

Set up governance framework and structures

This includes:

- The key functions and bodies that may be needed to enable the organisation to successfully deliver on its AI roadmap and strategy, and ensure the appropriate checks and balances, controls, decision-making, expertise, advice and guidance are there. More generally, this means ensuring all necessary stakeholders, from all relevant fields and including the C-suite, are involved and informed.
- The documents that reflect the organisation's AI roadmap, strategy and pledges externally, and those that enable the organisation to ensure implementation and compliance internally: AI principles, AI policies, dos and don'ts, etc.
- The organisation's risk management framework.
- The processes needed to ensure the safe and responsible development, uptake and deployment of AI.

3

Look at the interplay with, and how to leverage, what already exists within the organisation

This includes looking at frameworks, tools, principles, policies, etc. as may have been developed under data privacy or cybersecurity programmes, for instance.

4

Manage relationships with third parties

This may include such things as necessary auditing / due diligence, and contracting.

5

Monitor legal and market developments, and stay involved

This will be important, amongst other things, to ensure awareness, compliance and an adequate level of AI literacy within the organisation, as well as re potential engagement in the debate around helping shape norms.

AUTHORS

Dessislava Savova
Partner
Paris

Alexander Kennedy
Knowledge Director –
CE Tech Group, Paris

Rita Flakoll
Global Head of Tech Group
Knowledge, London

MORE PEOPLE WITH EXPERTISE IN THIS AREA

Zayed Al Jamil
Partner
London

Blanche Barbier
Associate
Paris

MORE PEOPLE WITH EXPERTISE IN THIS AREA

Jennifer Chimanga
Partner
London

Stella Cramer
Partner
Singapore

Juan Cuerva
Counsel
Barcelona

Caroline Dawson
Partner
London

Megan Gordon
Partner
Washington D.C.

Brian Harley
Counsel
Hong Kong

Ling Ho
Partner
Hong Kong

Arnav Joshi
Senior Associate
London

Jonathan Kewley
Partner
London

Violetta Kokolus
Partner
New York

Devika Kornbacher
Partner
Houston

Paul Landless
Partner
London

Holger Lutz
Partner
Frankfurt

Don McCombie
Partner
London

James McPhillips
Partner
Washington D.C.

Andrei Mikes
Counsel
Amsterdam

Claudia Milbradt
Partner
Dusseldorf

Josep Montefusco
Partner
Barcelona

Nadine Neumeier
Counsel
Frankfurt

Michihiro Nishi
Partner
Tokyo

Dieter Paemen
Partner
Brussels

Simon Persoff
Partner
London

Kate Scott
Partner
London

Gunnar Sachs
Partner
Dusseldorf

Phillip Souta
Global Director of
Tech Policy, London

Herbert Swaniker
Senior Associate
London

Jaap Tempelman
Senior Counsel
Amsterdam

Andrea Tuninetti-Ferrari
Counsel
Milan

Terry Yang
Partner
Hong Kong

Stavroula Vryna
Partner
London

CLIFFORD

CHANCE

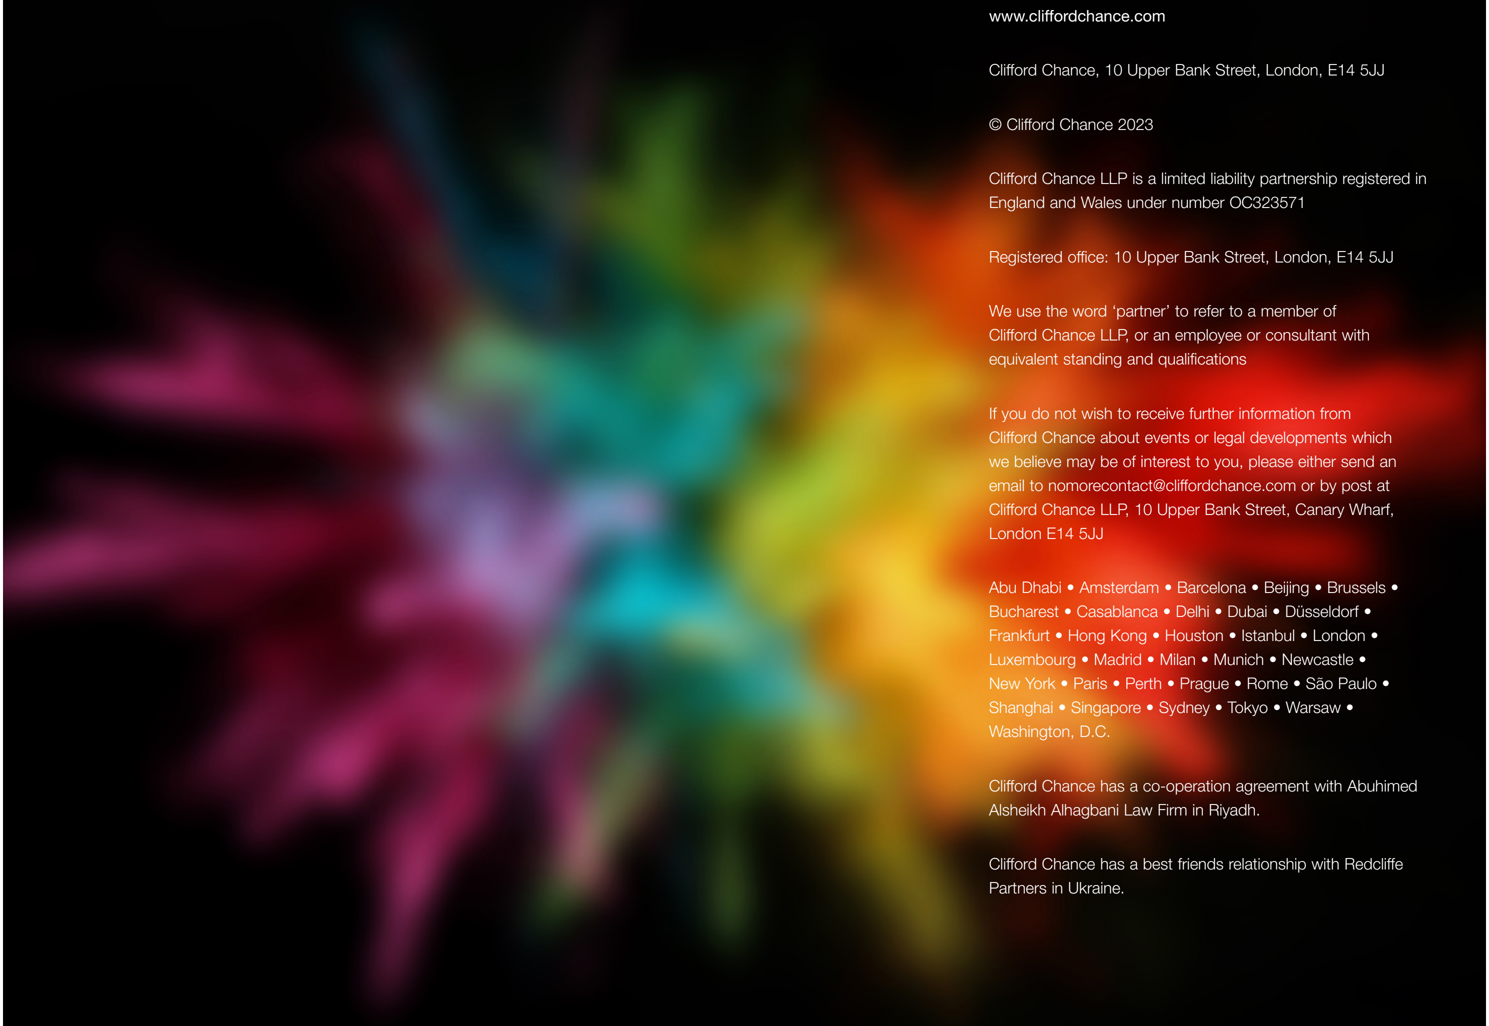

This publication does not necessarily deal with every important topic or cover every aspect of the topics with which it deals. It is not designed to provide legal or other advice.

www.cliffordchance.com

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

Clifford Chance LLP is a limited liability partnership registered in England and Wales under number OC323571

Registered office: 10 Upper Bank Street, London, E14 5JJ

We use 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 • Rome • São Paulo • Shanghai • Singapore • Sydney • Tokyo • Warsaw • Washington, D.C.

Clifford Chance has a co-operation agreement with Abuhimed Alsheikh Alhagbani Law Firm in Riyadh.

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