

Diritto intelligente

In this issue

- *[EDPB web scraping guidelines: Five open questions for the growth of AI](#)*
- *[Character.AI and the Italian DPA's fine: Europe's first test case on child protection in the age of generative AI](#)*
- *[DMA and interoperability for AI services: The commission requires Google to open up android to competing assistants](#)*
- *[AI Act transparency obligations: The top 10 questions after the EU guidelines](#)*
- *[AI Act investigations: How to handle the first information requests](#)*

Contents

EDPB web scraping guidelines: Five open questions for the growth of AI 4

Character.AI and the Italian DPA’s fine: Europe’s first test case on child protection in the age of generative AI 7

DMA and interoperability for AI services: The commission requires google to open up android to competing assistants 9

AI Act transparency obligations: The top ten questions after the EU guidelines 11

AI Act deepfakes and AI-generated text: What deployers must do from 2 August 2026 16

AI Act investigations: How to handle the first information requests 18

AI criminal liability in Italy: What boards must know now.....21

CADA: The European strategy for cloud and AI between industrial support, sovereignty and public procurement. 24



Editorial

AI accountability is becoming a boardroom issue

If one message runs through this issue of *Diritto Intelligente*, it is that AI regulation is no longer primarily about understanding what the rules will eventually say. It is about being able to demonstrate, today, who made the decision, on what basis, with what safeguards, and what evidence exists to prove it.

The shift is visible across the entire AI lifecycle. The EDPB's new web scraping guidelines reach back to the data used to train and fine-tune models, requiring organisations to confront questions of legal basis, transparency, minimisation, source selection and even the technical measures used to prevent problematic data from entering the dataset. At the other end of the lifecycle, the AI Act transparency rules now require businesses to explain when users are interacting with AI and to label certain AI-generated content. The common denominator is not the technology. It is accountability.


The enforcement picture makes this even more concrete. The first AI Act information requests are no longer hypothetical regulatory exercises. A company's response can itself become part of the case, which means that the first letter from an authority is not simply a request for information but the beginning of a defence. The Italian framework then takes the issue further, with the introduction of criminal liability for failures concerning the safety and human oversight of high-risk AI systems, together with corporate exposure under Legislative Decree 231/2001. The question for management is therefore changing from "What happens if we are fined?" to "Can we demonstrate why we approved this system in the first place?"

The same principle appears in the Garante's action concerning Character.AI and in the warning on workplace emotion inference. In both cases, compliance cannot be retrofitted after the technology has been deployed. Age safeguards, DPIAs, transparency, intended purpose and technical limitations have to be considered when the system is designed and before it reaches users. Good intentions do not replace governance, and a compliance policy that cannot be demonstrated in practice is unlikely to provide much protection when the regulator arrives.

There is, however, another dimension to this issue. Europe is not only regulating how AI is developed and used. It is increasingly trying to shape the market in which that AI develops. The Commission's intervention on Google's Android ecosystem and the proposed Cloud and AI Development Act show a European strategy that combines competition, interoperability, industrial policy, cloud sovereignty and public procurement. The difficult question is whether Europe can build an AI ecosystem that is both trusted and competitive, without regulation itself becoming an obstacle to scale.

This is why the AI compliance file is becoming much more than a legal document. It is becoming part of corporate governance, procurement, product development and, increasingly, the defence strategy of the company.

AI regulation is becoming operational. Accountability is becoming personal. And the organisations that will navigate this transition best will be those that can prove what they did before anyone asked them to.



Giulio Coraggio

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

EDPB web scraping guidelines: Five open questions for the growth of AI

Author: *Giulio Coraggio*

The EDPB web scraping guidelines finally set out how the GDPR applies to the data that trains generative AI – and they raise open questions that could determine whether AI models can still be built in Europe.

On 7 July 2026, the European Data Protection Board adopted [Guidelines 03/2026 on web scraping in the context of generative AI](#), opening them to [public consultation until 30 October 2026](#). They came out of the [same plenary](#) that produced the new anonymisation guidelines and the final blockchain guidance, which tells you something about the direction of travel: the EDPB is systematically mapping the GDPR onto every layer of the AI stack.

And this layer is the foundation. No scraped data, no training dataset. No training dataset, no model. Whatever the guidelines say about scraping, they say about the feasibility of developing artificial intelligence in the European Union.

The document is only 22 pages. Do not let that mislead you. The level of granularity expected from controllers is remarkable, and the practical consequences for anyone training or fine-tuning a model, not just the large language model developers, but any company that scrapes, directly or through a vendor, are significant.

What the EDPB web scraping guidelines actually say

Let me start with the uncontroversial part, because the useful debate is elsewhere.

The GDPR applies to web scraping whenever the operations involve personal data: extraction, cleaning, structuring, storage. “Publicly available” is not a magic exemption, and the EDPB repeats it with no ambiguity. The guidelines cover two scenarios – an organisation that scrapes itself or instructs a third party to do so, and an organisation that reuses a dataset scraped by someone else. Pure data brokers who scrape and resell without training models are left out, as is public sector scraping.

From there, the guidelines walk through the principles:

- **Roles.** The scraper is not automatically the controller. Where a scraper builds a dataset on the documented instructions of an AI developer, it may be a processor, with the developer as controller. Where a developer reuses a dataset already scraped, each party answers for its own processing, and the original scraper is not, in principle, responsible for the reuse. Joint controllership arises where collection criteria are set jointly.

- **Transparency.** Individual privacy information notice is often impossible or disproportionate, and Article 14(5)(b) GDPR can be relied on but only after a real balancing exercise on the dataset as a whole, weighing the number of data subjects, the age of the data and the safeguards adopted. The fallback is not silence: the controller must publish the privacy information notice, including a precise indication of the sources, ideally with domain names, URLs in searchable format and collection dates.
- **Data minimisation.** A long list of expected measures, before and after collection: consider synthetic data first, define precise collection criteria, run a data mapping exercise, filter out categories that are not needed, exclude websites structurally likely to contain sensitive data or data of minors, exclude sites that oppose scraping via robots.txt, ai.txt or CAPTCHA, apply syntax-based filters during collection, and anonymise or pseudonymise where feasible.
- **Accuracy.** Scrape from reliable, maintained sources, timestamp the data, validate before training. And, this is easy to miss, accuracy also bites on the output: if the model is expected to generate personal data once on the market, that output must be accurate too.
- **Legal basis.** Consent is off the table for indiscriminate collection, and the EDPB is explicit that publishing data online is not consent to scrape it. Legitimate interest under Article 6(1)(f) remains the realistic route, subject to the three cumulative conditions of the EDPB Guidelines 1/2024 and [Opinion 28/2024 on AI models](#): a real, lawful, precisely articulated interest; necessity, meaning no equally effective and less intrusive alternative; and a balancing test.
- **Special categories.** Prohibited in principle. Scraping that captures Article 9 data needs both an Article 6 basis and an Article 9(2) derogation. The EDPB accepts that the CJEU’s reasoning in *GC & Others* (C-136/17) developed for search engines may apply to incidental and residual collection, within the controller’s “responsibilities, powers and capabilities”, provided technical and organisational measures prevent collection and dissemination. But there is no blanket exemption and no blanket assessment: case by case, every time.

On balance, the EDPB web scraping guidelines are a serious, well-constructed document, and considerably more pragmatic than some feared. The recognition that residual sensitive data does not automatically make training unlawful matters enormously. So does the acknowledgement that data minimisation does not prohibit training on large volumes of data.

But improvements and difficulties can coexist. And a consultation is exactly the moment to say where the difficulties are.

Why the EDPB web scraping guidelines matter for the growth of AI

Here is the strategic point. Europe has spent the last two years debating whether it can host AI development or merely consume it. The Digital Omnibus, the competitiveness agenda, the Draghi report – the whole conversation assumes that European companies can actually build models here.

The training dataset is where that assumption is tested. If the compliance cost of assembling a lawful dataset is high but predictable, European AI development survives and the guidelines become a genuine competitive asset, because trustworthy data provenance is worth something. If the cost is unquantifiable, the rational decision for any board is to buy a model trained elsewhere and deploy it which achieves nothing for the protection of data subjects, since the same data was scraped anyway, only by someone outside the reach of the EDPB.

This is a point worth raising constructively in the consultation, because it is one the EDPB is well placed to address: guidance of this kind is naturally most visible to the organisations already committed to complying with it.

There is also a structural point that deserves attention. The measures the EDPB expects, precise collection criteria, source exclusion, filtering, mapping, pseudonymisation, memorisation controls, are not merely documentary. They are engineering. Larger developers are generally equipped to absorb them. A European scale-up fine-tuning an open-weight model for a vertical use case may find the same expectations proportionately heavier. Some calibration of the guidance to the size and scope of the processing would help ensure that the framework supports innovation across the market, and not only at the top of it.

What to do now: a practical checklist

Before the debate on the open points, there is work that can start immediately, because none of it depends on how the consultation ends. If your organisation scrapes, commissions scraping, or buys scraped datasets:

- 1. Map who you are.** Controller, joint controller or processor, per processing activity. If you instruct a vendor, check the instructions are documented and specific on sources and categories – that is what makes them a processor rather than a joint controller.
- 2. Re-open your legitimate interest assessment.** A generic LIA will not survive. Articulate the interest precisely (commercial or research, internal or external), document why less intrusive alternatives – synthetic or pseudonymised data, narrower criteria – were not equally effective, and record the balancing test with the mitigating measures actually implemented.

- 3. Write down your collection criteria before you crawl.** Untargeted crawling is the hardest position to defend, because the EDPB links it directly to the controller not knowing what it holds.
- 4. Honour opposition signals.** Robots.txt, ai.txt, CAPTCHA and authentication walls. Log the exclusions – the evidence matters more than the practice.
- 5. Exclude structurally risky sources.** Sites used mainly by minors, and sources structurally likely to carry financial, location or Article 9 data.
- 6. Timestamp and validate.** Record collection dates, prefer official and maintained sources, spot-check samples before training.
- 7. Publish an Article 14(5)(b) notice that actually works.** Categories of data, purposes, legal basis, crawler characteristics, and the source list – domain names and URLs in searchable format where possible, with the reasons for any omissions.
- 8. Address the model, not just the dataset.** Memorisation and regurgitation controls, output filtering and accuracy of generated personal data all feed back into the balancing test.
- 9. Paper the supply chain.** For purchased datasets: sources, exclusion criteria, audit rights, warranties on how the dataset was assembled, and a contact point for the originating controller.
- 10. Run a DPIA – and consider publishing it.** The EDPB lists publication among the appropriate safeguards where individual notice is disproportionate.

Five open questions before 30 October 2026

- 1. How wide is the “responsibilities, powers and capabilities” door?** *GC & Others* was about a search engine responding to referencing requests. Extending it to the untargeted crawling of the open web is a real analytical leap. If it works, incidental sensitive data is manageable. If a data protection supervisory authority reads it narrowly at enforcement stage, most large-scale scraping is unlawful under Article 9 no matter what filters were applied. Which reading will prevail?
- 2. Does robots.txt now have legal effect?** The guidelines treat opposition signals, robots.txt, ai.txt, CAPTCHA, both as a data minimisation expectation and as a factor in reasonable expectations, with Examples four and five turning on precisely that. A technical convention that carries no legal status of its own is therefore acquiring real weight in a GDPR balancing test. That is a meaningful development and one worth examining openly during the consultation. What happens where the signal is ambiguous, changed after collection, or set by a platform on terms its own users have not chosen?

- 3. What is the temporal scope?** Models already trained cannot easily unlearn, the EDPB says so itself when discussing the impossibility of deleting personal data from a trained model. Do the EDPB web scraping guidelines apply to datasets assembled before July 2026? To fine-tuning a model whose base training predates them? The document is silent, and silence here is expensive.
- 4. How does this interact with the AI Act and the Digital Omnibus?** The AI Act has its own data governance requirements for training data. The Digital Omnibus reopened the definition of personal data before the contested amendment was dropped following the EDPB and EDPS Joint Opinion 2/2026. Controllers now face parallel regimes, drafted by different institutions, on the same dataset. Who reconciles them, and when?
- 5. Is the source list workable?** Publishing domain names and URLs in searchable format, with collection periods, is transparency in its purest form. It is also, for an untargeted crawl across millions of URLs, a substantial disclosure of what may be a competitive asset. Is there a proportionate middle ground between a meaningless list of source categories and a full crawl inventory?

My view


The EDPB web scraping guidelines are not the obstacle some will claim. They are closer to a roadmap than a wall, and the flexibility on transparency and on residual sensitive data shows an institution that has listened. The direction is sound. What remains is uncertainty at exactly the points where a general counsel has to give a clear answer to a business that wants to invest.

That uncertainty is addressable, and the consultation is the right forum in which to address it, before divergent national practice develops around it.

So I will put the same question to you that I have been putting to colleagues over the last week: which of these open questions is the one your organisation cannot live with? I suspect the answers differ enormously depending on whether you train, fine-tune or merely deploy, and that difference is itself something the EDPB should hear before 30 October 2026.

On a related topic, you may find of interest the article [“EDPB Anonymisation Guidelines: What They Mean for AI Systems”](#) on this blog, which covers the companion guidance adopted at the same plenary.

The implications of the EDPB web scraping guidelines depend heavily on whether your company trains, fine-tunes or deploys AI models, and on how your training data was assembled.



Character.AI and the Italian DPA's fine: Europe's first test case on child protection in the age of generative AI

Author: Enila Elezi

On 3 July 2026, Italy's Garante per la protezione dei dati personali fined Character Technologies, Inc. for violations relating to the protection of minors, age verification, and data protection impact assessment.

Character.AI belongs to the category of so-called "AI companion" apps: services that allow users to converse with virtual characters generated by language models, often with significant emotional involvement. By now, these services attract teenagers is well documented; what is new is that authorities have moved from warnings to sanctions.

The service and the investigation

Character.AI is a generative artificial intelligence service developed by Character Technologies, Inc. (Menlo Park, California) that allows users to create and chat with virtual "Characters", either pre-existing on the platform or created by users themselves, through proprietary language models. The service launched in beta in September 2022 on the web and in May 2023 as a mobile application. The final version has been available since 8 April 2024 at character.ai, with an Italian-language interface and a privacy notice addressed to users in the European Economic Area.

The investigation was opened ex officio in November 2024 and concluded with the decision of 3 July 2026, in which the Garante found five distinct GDPR violations.

The five violations

- 1. Inadequate privacy notice (Articles 12(1), 13(1)-(2), 14(1)-(2) GDPR).** The privacy notices of 23 October 2023 and 27 August 2025 contained gaps regarding processing purposes, legal bases, retention periods, and international transfers. The first version was available only in English, despite the service being offered in Italian; the subsequent translation contained inaccuracies and inconsistencies between data categories, purposes, and legal bases.
- 2. Inadequate information on model training (Article 14(1)-(2) GDPR).** Character did not adequately inform data subjects residing in Italy, including non-users, of the possible use of data collected from the web (web scraping) to train its models. The Garante did not, however, find a violation of Article 21(1) and (4) on the right to object, noting that EDPB Opinion 28/2024, which clarified the applicable standards, post-dated the conduct in question.

- 3. Inadequate technical and organisational measures for the protection of minors (Articles 24(1), 25(2) GDPR).** A technical verification conducted by the Garante's Office demonstrated that it was possible to register by declaring an age of 15, despite the threshold being set at 16 for the EEA. The self-declared minor's profile was public by default. These shortcomings were assessed in light of the elevated risk arising from the vulnerability of the users involved, the sensitive nature of conversations with the chatbot, and press reports of self-harm episodes linked to similar services.
- 4. Late data protection impact assessment (Articles 5(2), 35(1) GDPR).** The DPIA should have preceded the launch of the service: the technology employed (generative AI) was alone sufficient to classify the processing as high-risk, together with large-scale processing and the involvement of minors. The first DPIA dates from November 2024, more than two years after launch, and was updated in September 2025. The initial version lacked an analysis of the specific risks to minors and to model training.
- 5. Late designation of the EU representative (Article 27(1) GDPR).** VeraSafe Ireland Ltd. was designated only on 31 May 2025, although the service was directed at EEA users and available in Italy since April 2024. A separate allegation concerning a non-functioning contact link (Article 27(4)) was not upheld, having been considered a clerical error that was promptly corrected.

The fine and corrective measures

The fine amounts to EUR158,000, calculated according to EDPB Guidelines 04/2022 and treating the violations as linked processing operations under Article 83(3) GDPR. The ceiling was determined by the most serious violation, the transparency breach, subject to the maximum under Article 83(5).

The gravity of the violations was judged "medium." The Garante found negligence, not wilful misconduct, and not of a serious degree, and recognised several mitigating factors: full cooperation during the investigation, absence of prior violations, progressive improvement of compliance (with the launch of a dedicated experience for minors, Parental Insight tools, and, from November 2025, the removal of open-ended conversations for under-18s), as well as the company's limited size and resources as an operator in a competitive market.

The corrective measures, to be implemented within 120 days, include: bringing the privacy notice into compliance regarding retention periods and international transfers; verifying the lawfulness of data used for training, with deletion in the absence of a compatible purpose; ensuring that the age filter effectively blocks EEA users under 16; preventing repeated registration attempts by minors already blocked; setting minors' profiles to private by default. Character must report to the Garante on compliance within the same timeframe.

Accountability and foreseeability of risk

The decision should not be read as a case about age verification technology as such. The Garante makes this explicit: the GDPR does not impose a single technical standard, and the lack of a harmonised standard does not, in itself, constitute a violation.

The central point lies elsewhere. Anyone deploying generative artificial intelligence directed at minors, or that foreseeably attracts them, must anticipate the risks inherent in the technology and in its user base from the design stage. Safeguards cannot arrive afterwards, under pressure from the regulator or the media.

The decision confirms that the timing of a DPIA is assessed based on when risk indicators already existed, not based on when the company decided to formalise its internal analyses. An innovative technology such as generative AI was in itself sufficient to trigger the impact assessment obligation before launch.

This is a common pattern among technology start-ups: first launch the product and acquire users, then document compliance when resources arrive or the regulator's attention does. The Garante states that this approach is incompatible with the accountability principle under the GDPR.

Relevance for businesses and digital platforms

The implications extend beyond the Character.AI case. Any company offering generative AI services or interactive digital services capable of reasonably attracting minors, even without explicitly targeting them, should expect supervisory authorities to verify:

- whether privacy notices are genuinely accessible and linguistically accurate in the language of the target market;
- whether model training practices are transparently communicated, including to non-users;
- whether age verification measures actually work and are not merely documented on paper;
- whether DPIAs are conducted before launch and not retrofitted afterwards;
- whether the EU representative is designated and reachable from the first day of operations in the EEA.

Companies developing or offering generative AI services, chatbots, or digital assistants, especially those attractive to younger users, therefore have every interest in acting preventively.


Privacy notices should be checked for accuracy and completeness in every language of the markets served: a formally correct translation is not enough if the technical-legal content contains inaccuracies. The data supply chain used for training must be mapped and documented with concrete evidence, not generic assurances, indicating whether and how personal data, including data collected from public sources, feeds the training and refinement of models, with effective opt-out mechanisms.

Age verification mechanisms should be fully tested, verifying actual operation rather than design intentions. Every account attributable to a minor should be set to private by default with limited functionality. DPIAs should be conducted and documented before launch whenever generative AI, large-scale processing, or vulnerable users are involved: "innovative technology" alone is a sufficient trigger.

The EU representative should be designated and made reachable from the moment of entry into the European market. Cooperation with the authority and demonstrable progressive improvement should be considered mitigating factors capable of affecting the outcome of sanctions: the Garante expressly credited them in reducing the fine.

Conclusion

The decision sets a precedent for how European authorities intend to address generative AI services with a user base predominantly composed of minors: through rigorous oversight of risk-based accountability, from design to deployment, without imposing prescriptive technical standards.



DMA and interoperability for AI services: The commission requires google to open up android to competing assistants

Author: Josaphat Manzoni

On 16 July 2026, the European Commission adopted a specification decision against Google pursuant to Article 8(2) of the Digital Markets Act (DMA), following proceedings initiated on 27 January 2026. The decision identifies the concrete measures that Google will have to implement to ensure effective interoperability of 11 functionalities of the Android operating system relevant to AI services, in compliance with Article 6(7) of the DMA.

Article 6(7) of the DMA requires gatekeepers to provide third-party developers with free and effective interoperability with the hardware and software functionalities controlled through their operating system. In the case of Google, this concerns Android, which constitutes the operating system of approximately 60% of mobile devices used in Europe. According to the Commission, Google currently grants its own AI assistant, Gemini, a level of deep access to key functionalities of the operating system that is not available to competing AI assistants, effectively preventing alternative providers from offering comparable assistive experiences.

Rationale and expected impact of the decision

AI assistants are becoming one of the primary means of interaction between users and their mobile devices. Gemini, supported by Google's entire technology stack, is in a privileged position to become the dominant AI offering on mobile devices. Android's distribution network enables Google to deliver its own AI services directly to users, leveraging an integration with the operating system that competitors cannot replicate.

The Commission considered that effective competition and genuine innovation in the AI services market can only develop if alternative providers have a level of access to the Android operating system analogous to that enjoyed by Google's services. The measures adopted are therefore intended to enable competing AI assistants to offer equally functional and performant experiences, without entailing compromises in terms of security or performance.

The expected benefits also extend to developers whose apps can be used by an AI service, such as email or messaging apps, who will benefit from greater visibility and increased interaction through integration with alternative AI assistants.

The measures Google must implement

The decision covers 11 Android functionalities, grouped into four macro-categories representing the fundamental elements on which AI services are built.

Invocation. The first category concerns the ways in which the user can initiate interaction with an AI service. The measures provide that third-party services may be activated via a long-press of the home button or navigation handle, an access point currently reserved for Google's services. In addition, users must be able to activate an AI assistant via a voice command, a functionality currently reserved for Google's services. The user will thus be able to invoke their preferred AI assistant at any time, even with the screen off or in power-saving mode.

Context. The second category concerns the ability of AI services to understand the user's needs and anticipate their requirements. The measures guarantee third-party services centralised access to apps' data stored on-device, currently available only to Google's services. Access to context-aware intelligence is also provided, namely the ability to offer proactive personalised suggestions based on the user's context, without the user having to activate the assistant. Finally, third-party services will have access to ambient data, i.e. the continuous real-time data stream from the device's sensors, under the same conditions as those applicable to Google's services in terms of user consent.

Actions on apps and the OS. The third category concerns the ability of AI services to perform tasks on behalf of the user. The measures provide access to structured on-device integration, which enables AI services to interact with other installed apps and perform actions such as sending a message, creating a note, or scheduling a meeting, including specific functionalities of certain Google apps. The measures also guarantee access to the screen automation functionality, which enables AI assistants to automate complex tasks within other apps, imitating user behaviour in a separate virtual window running in the background. Finally, system integration enables AI services to interact directly with device settings and system functionalities.

Access to resources. The fourth category concerns the hardware and software resources necessary for task execution. Third-party services will be able to access system-level on-device models already pre-installed on Android devices, with guarantees of equal access in terms of performance compared to Google's services. The measures also enable third-party providers to install, run, and use their own models directly on-device under the same conditions enjoyed by Google's models, both in terms of access to hardware resources and background execution. Finally, the decision guarantees equal access to the background execution functionality, which allows apps to perform actions in a timely manner even when the user is not actively using the application.

With respect to all the functionalities described above, Google must ensure interoperability free of charge, on all Android devices in the ecosystem, including those of third-party manufacturers. The interoperability solutions must be equally effective compared to those available to Google's services and may not impose unnecessarily complex settings, additional friction for the user, or be conditional upon the AI provider holding a default role. Google must also provide third parties with complete documentation, testing tools, and technical assistance, as well as make any new functionalities available at the same time as they become accessible to its own services.

Safeguards for system integrity and restricted features

The decision recognises Google's ability to adopt measures that are strictly necessary and proportionate to safeguard the integrity, security, and privacy of its operating system. Such measures must, however, be based on transparent, objective, precise, and non-discriminatory conditions, applicable equally to Google's services and those of competitors. In addition, the decision draws a distinction between open-access features and restricted features. For the latter, which include among others centralised access to apps' data, context-aware intelligence, structured on-device integration, and screen automation, Google may require that the developer demonstrates compliance with certain requirements before granting access.

To this end, the decision requires Google to establish a Qualified AI Assistant Programme to verify that AI assistants meet the prescribed requirements. The eligibility

conditions may concern, among other things, the assistant's advanced reasoning capabilities, respect for user intent, protection against the main risks associated with the use of AI agents, privacy, security, and transparency requirements, as well as developer reliability. Certification may be obtained directly from Google or through independent Trusted Certification Authorities. Google must conclude the assessment of each application within four weeks of receipt of the request and must publish the final conditions of the programme by 1 May 2027.


Such instruments, while functional to ensuring the quality and security of services offered on AI assistants, could nevertheless be used by Google to limit in practice the pro-competitive impact of the decision. The definition of eligibility conditions, certification procedures, and access conditions to the most sensitive functionalities represent levers that, if applied disproportionately, could unduly restrict competitors' access. It will therefore be the Commission's task to carefully monitor that such mechanisms are not used abusively by Google.

Conclusions

The Commission's specification decision represents a significant intervention aimed at ensuring that competition in the AI assistants market on Android is not hindered by the control exercised by Google over the key functionalities of its operating system. If effectively implemented, the measures will enable alternative providers to offer assistive experiences comparable to those of Gemini, contributing to a more contestable market to the benefit of both developers and consumers.

The measures must be implemented with the release of Android 18, by 1 August 2027, with the exception of concurrent hotword detection, scheduled for Android 19 by 1 August 2028.

It should moreover be noted that the DMA is already fully applicable and that Google is already required to ensure interoperability with the functionalities of its operating system. The granting of an implementation period until August 2027 signals that the Commission considered the required measures sufficiently complex to justify a reasonable timeframe for their implementation, balancing the need for effective enforcement of the regulation with the necessity of ensuring technically sound and functional solutions.



AI Act transparency obligations: The top ten questions after the EU guidelines

Author: *Giulio Coraggio*

The **AI Act transparency obligations** finally have their rulebook: on 20 July 2026, the European Commission published the [final guidelines on the implementation of Article 50 of the AI Act](#). Thirteen days before 2 August 2026 when the relevant provisions of the EU AI Act become applicable.

The timing tells you everything about how this file has been managed in Brussels. But it also removes the last excuse for waiting. After months of projects on this topic, my impression is that the **AI Act transparency obligations** under Article 50 are the ones destined to touch the highest number of companies – far more than currently realise it. This is not a rule for those who build high-risk AI systems. It is a rule for anyone who puts artificial intelligence in front of people, through a chatbot, an image, an audio file, a video or a published text.

The guidelines are non-binding, and only the Court of Justice can give an authoritative interpretation of the AI Act. In practice, though, they will be the primary reference used by market surveillance authorities across all 27 Member States when they assess Article 50 compliance. That makes them the closest thing to a rulebook that businesses will have on 2 August 2026.

Below are the ten questions I am asked most often, answered in the light of the new guidelines, the Code of Practice on Transparency of AI-Generated Content, the Digital Omnibus and the Italian implementing decrees under Law No. 132/2025.

What the guidelines add to the AI Act transparency obligations

Three documents now sit on the same table, and they do different jobs.

The **guidelines** define the scope: who is caught, what counts as an interaction, what counts as a deep fake, when an exception genuinely applies. The [Code of Practice on Transparency of AI-Generated Content](#), published on 10 June 2026 and confirmed by the Commission and the AI Board as an adequate voluntary tool, provides the practical framework for marking and labelling. The [EU set of icons](#) gives deployers a ready-made visual label.

The Code is voluntary. The obligations are not. And that distinction is where most of the compliance strategy now sits, as I explain at question nine.

1. Which AI Systems Are Actually Caught – and Who Is About to Find Out

Article 50 builds four distinct obligations, and they are worth keeping separate because they fall on different actors:

- **Systems interacting directly with people** (Article 50(1)): chatbots, voice assistants, automated customer service. The provider must design the system so that users know they are talking to a machine.
- **Generative AI systems** (Article 50(2)): whoever produces synthetic text, images, audio or video must mark the outputs in a machine-readable format and ensure they are detectable as artificial.
- **Emotion recognition and biometric categorisation** (Article 50(3)): the deployer must inform the people exposed to the system.
- **Deep fakes and text on matters of public interest** (Article 50(4)): the deployer must disclose the artificial nature of the content.

The point I always flag is this: the addressees are not “AI producers”, but the far broader category of businesses using AI in their daily operations. The e-commerce with a support chatbot. The communication agency generating visuals and copy. The company publishing press releases or editorial content built with AI. The retailer testing customer analytics.

These are organisations that often do not perceive themselves as “AI Act subjects” at all – and that are squarely within scope.

The guidelines also confirm something many had hoped to avoid. **AI agents are covered**, and they must disclose both their artificial nature *and* the person on whose behalf they are acting. Where a provider cannot determine in advance whether an agent will interact with a human, the agent must be designed to disclose itself in every situation where such interaction is reasonably likely.

In my experience, the most frequent surprise concerns SMEs that adopted off-the-shelf generative tools without asking what obligations followed, on the assumption that everyone is doing it, so it will be fine.

It will not be fine.

2. Provider or Deployer? The qualification is not carved in stone

In the vast majority of cases, a company that integrates a third-party chatbot into its website is a **deployer**: it uses the system under its own authority, but did not build it.

Be careful, though, because the qualification can shift. If a business substantially modifies the system, places it on the market under its own name or trademark, or changes its intended purpose, it can slide into the **provider** position and inherit far heavier obligations.

The guidelines are explicit on this point, and it happens more often than people think.

The split of duties works as follows:

- The **provider** must design the system so that the disclosure of its artificial nature is technically possible and effective (Article 50(1)), and must mark generative outputs (Article 50(2)).
- The **deployer** must ensure that the disclosure actually reaches the user in the concrete context of use, and carries the labelling obligation for deep fakes and public-interest texts (Article 50(4)).

Two clarifications in the guidelines deserve attention. First, “authority” over an AI system does not require technical control – it is about deciding whether and how the system is used. Second, individual employees acting under the instructions of a company are not separate deployers: the legal person is.

The correct reading is a chain. The supplier provides the technical capability; the company using it activates and displays it. If either link breaks, the obligation is not met – which is exactly why the contract with the supplier becomes decisive, as I cover at question eight.

3. What a compliant AI disclosure looks like – and what is not enough

Here the guidelines are genuinely useful, above all in telling us what does **not** work. The following are expressly listed as insufficient when used alone:

- disclosures contained only in terms and conditions, URLs or documentation;
- machine-readable markings, such as metadata or watermarks, that the user cannot perceive at the point of interaction;
- unclear or ambiguous signals, such as a generic reference to an “assistant”, or human-like representations that may mislead;
- generalised statements such as “services on this website use AI”;
- purely technical descriptions such as “this system uses LLMs”, without explaining what that means for the user.

The wording that works

What is required instead is a notice that is **clear, in plain language, and placed at the moment the user comes into contact with the system**. In practice I suggest a direct formulation – along the lines of “You are interacting with a virtual assistant based on artificial intelligence” – visible before or at the start of the conversation, not hidden behind a click. The guidelines reward a combination of formats: a plain textual message accompanied, where useful, by persistent visual elements.

The guiding test I keep in mind is simple. A reasonable user, at the point of contact, must understand immediately and without effort that there is no person on the other side.

When the interaction is genuinely “obvious”

And the **“obvious interaction” exception is narrower than most people assume**. The guidelines state that it should be limited to cases where there is almost no doubt left about the nature of the interaction. Code assistants used only by professional developers, internal employee-facing tools for trained staff, or diagnostic support tools used only by health professionals can qualify. A customer-facing helpdesk chatbot on a public website does not.

One more practical point: in riskier contexts – AI companions, financial or health advice, complaints handling, or interactions with children and elderly users – a single notice at the start may not be enough. Periodic reminders are likely to be necessary.

4. What counts as a deep fake, and why your intentions are irrelevant

The Commission has adopted a deliberately broad interpretation, and this is a point on which I invite clients not to deceive themselves.

The definition captures AI-generated or manipulated image, audio or video content that appreciably resembles existing persons, objects, places, entities or events, and that would falsely appear to a person to be authentic or truthful. It does not require a real person to be depicted, and it does not require a sophisticated manipulation.

And here the answer is blunt: **yes, the obligation applies even in the complete absence of any intention to deceive**. The guidelines confirm that the assessment is objective. Transparency is not a function of your good faith.

That is the most widespread error I see: thinking “it is obviously fake, I am not trying to fool anyone” and concluding that no label is needed.

The audience you must have in mind

Two further points from the guidelines are worth internalising. Unlike the Article 50(1) test, the deep fake assessment is **not** based on a hypothetical average person: it must take into account the diverse composition of the reasonably foreseeable audience, including children, the elderly and people with lower digital literacy. Content that would mislead that part of the audience qualifies.

Conversely, content that defies the laws of nature – dragons, flying humans, elephants driving cars – falls outside the scope, because it has no potential to mislead.

Commercially, the most relevant example is advertising. An AI-generated image of a product that makes it appear more appealing, of better quality or different from the real thing is a deep fake. A real car photographed against an AI-generated background, where the ad does not mislead about the product itself, is not.

There is a lighter regime for content that is **evidently** artistic, creative, satirical or fictional: there, disclosure must simply be appropriate and must not hamper the enjoyment of the work. But a full exemption does not exist. And where content mixes an informative and a creative character, the guidelines say the informative character always prevails.

5. AI-Generated text on matters of public interest: Who labels it, and how

The obligation falls on the **deployer**, meaning the party publishing or disseminating AI-generated or manipulated text on matters of public interest. Think of an AI-generated summary of an article on a newspaper website, AI-manipulated corporate reports containing investor information, or a weather warning published by a public institute.

There is, however, a significant exception: where the content has undergone **human review or editorial control** and a natural or legal person holds **editorial responsibility**, no label is required. Both conditions must be met.

The guidelines close several doors here that were left ajar in the draft:

- superficial or purely formal checks – spell-checking, grammar correction, the mere existence of an editorial policy, cursory sign-off – do not qualify;
- fact-checking the accuracy of the content is described as a **minimum** requirement of the review;
- if AI is used to modify or reformulate the content **after** editorial sign-off, the exception becomes void;
- the identity and contact details of the person or function holding editorial responsibility should be publicly available in an easily findable location.

On the “how”, the EU has now published a set of icons with three variants: a **basic icon**, a **Fully AI-Generated** icon, and a **Partially AI-Modified** icon. They are free to use, and user testing showed that performance improves when the icon is accompanied by a short text label. Their use is optional; the labelling obligation is not.

The operational advice I give to communication and marketing teams is to define an internal editorial policy now: when AI is used, who verifies, who takes responsibility, and when the label must be applied. A clear process beats case-by-case decisions every time.

6. Machine-Readable marking and the 2 December 2026 window

This extension comes from the Digital Omnibus package, and it needs to be read precisely, because it is narrower than it is often described.

The rule is that generative AI systems **placed on the market or put into service before 2 August 2026** benefit from a transitional period and must comply with the Article 50(2) marking obligation **by 2 December 2026**. Systems placed on the market from 2 August 2026 onwards get no tolerance at all: they must mark their outputs from day one.

The decisive criterion is therefore the **date of placing on the market or putting into service** of each individual system. Which is precisely where this intertwines with mapping: without an inventory recording, system by system, when each was adopted, you cannot even establish which deadline applies to you.

Three warnings I repeat constantly:

1. The extension covers **only** Article 50(2). Every other Article 50 obligation – chatbot disclosure, deep fake labelling, public-interest text labelling – applies from 2 August 2026.
2. Systems that are partly interactive and partly generative benefit only in relation to the marking obligation. The interaction disclosure is due on 2 August.
3. This is not permission to slow down. It is a handful of extra weeks to finish a technical implementation.

The quality requirements behind the marking

A caveat on timing. The amending regulation had been adopted by the EU legislature but was still awaiting publication in the Official Journal when the guidelines were issued, and it enters into force on the third day following publication. Until then, the December date operates as a planning horizon rather than as applicable law. I would build the plan around it, but not bet the compliance programme on it.

On the substance, the technical solutions must be effective, interoperable, robust and reliable, in line with the generally acknowledged state of the art. And the guidelines make an important point: **technical feasibility is an objective notion, not dependent on the resources of the individual provider**. You cannot argue that marking was unfeasible for you because your team is small.

There are limited carve-outs – genuinely industrial or business-to-business applications meeting three cumulative conditions, generative outputs confined within closed physical products, and real-time ephemeral content that is not stored or disseminated. They are narrow, and they are not a general escape route.

7. Emotion recognition and biometric categorisation: check the prohibition first

Here I have to reverse the question, because this is where I see the most confusion.

Before transparency comes the prohibition. **Emotion recognition in the workplace and in educational institutions is banned** under Article 5 of the AI Act, save for narrow medical or safety exceptions. So in a recruitment scenario, the issue is not “how do I inform the candidate”: it is that, as a rule, that system **cannot be used at all**. That is a distinction to clarify with HR teams immediately, because the exposure is not an information gap but a prohibited practice, carrying the highest penalty tier in the entire Regulation.

Where use is permitted – certain retail applications, for instance – Article 50(3) applies and the deployer must inform the people exposed to the system's operation. The guidelines confirm that this applies whether the system runs in real time or ex post, and that the information must reach everyone exposed, including children. Practical examples given include a pop-up before a game is launched, or a visible notice at each entrance to a room where visitors' facial images are captured.

On top of that sits the entire GDPR framework, because we are talking about biometric data and potentially special categories: legal basis, impact assessment, minimisation.

One technical point that matters enormously in practice. Many tools marketed as “sentiment analysis” simply analyse keywords in a conversation to understand its tone, without any biometric processing that would trigger the prohibition. That distinction was at the heart of the Italian Data Protection Authority's warning in the Myndoor case, which I analysed in [this article on AI sentiment analysis in the workplace](#). Before reasoning about the notice, verify how the technology actually works – and whether it is lawful at all.

8. The five mistakes i see most – and the clauses to renegotiate now

The recurring errors, in my experience, are five:

- burying the disclosure in the terms and conditions or the privacy policy, instead of placing it at the point of contact;
- relying on the technical watermark alone, forgetting the notice the user can actually see;
- using generic or euphemistic wording – “assistant”, “smart service” – that never declares the artificial nature;
- failing to map the systems, and therefore not knowing which obligations genuinely apply;
- treating compliance as a one-off project, with no ongoing ownership.

On the contractual side, I invite clients to renegotiate – or at least review – the following with their AI suppliers now:

- a **clear allocation of the provider/deployer roles** and of the related responsibilities;
- **warranties** that outputs are marked and detectable as artificial in line with Article 50(2);
- a commitment to make the **disclosure functionalities** and the necessary technical documentation available;
- **indemnities** for non-compliance attributable to the supplier;
- an obligation to **cooperate in the event of an authority inspection**;
- the **flow-down** of the obligations along the supply chain.

These clauses are frequently missing altogether from vendors' standard terms. That is not a drafting oversight you want to discover during an inspection.

9. Who enforces, with what powers, and what the fines look like

In Italy, the picture was settled by the implementing decrees under Law No. 132/2025, which I covered in detail in [this article on Italy's AI Act implementation](#). The pivot of operational supervision – including inspections and sanctions – is **ACN**, the National Cybersecurity Agency, while **AgID** is the notifying authority. The competences of sectoral regulators (Banca d'Italia, CONSOB, IVASS) and of the Garante remain intact within their respective areas.

On concrete powers, Article 50 plugs into the European market surveillance system. In practice that means requests for documentation and information, on-site inspections, the power to order corrective measures and, in the most serious cases, the power to require withdrawal or recall of a non-compliant system.

Penalties reach **EUR15 million or 3% of total worldwide annual turnover, whichever is higher**.

For SMEs and start-ups, the lower of the two applies.

There is no exemption based on size: an individual professional or a micro-enterprise publishing AI-generated content without the required disclosures is fully within the perimeter.

The Italian decree adopts a graduated and proportionate approach, calibrating penalties to the degree of responsibility along the supply chain. That is a proportionality criterion, not a safe conduct.

Signing the code, or proving it yourself

And here is the point I consider most strategically relevant in the guidelines. Signatories of the Code of Practice will see supervisory activity focused on whether they have implemented the measures it contains. Non-signatories are expected to demonstrate compliance by other adequate means, to **explain** how their measures achieve compliance, and – in the Commission’s own words – to carry out a **gap analysis** comparing what they have implemented against the Code. They should also expect a larger number of requests for information and requests for access.

In other words, not signing is a legitimate choice. It is also a choice to demonstrate compliance with the AI Act transparency obligations the hard way, carrying a heavier evidentiary burden. The deadline to be included in the first published list of signatories fell on 22 July 2026, although joining later remains possible.

One final observation from enforcement practice: as we have seen repeatedly in the Garante’s decisions, adopting corrective measures after the fact is generally not considered sufficient to avoid a sanction. Prepare the documentation **before** the inspection, not after.

10. The three priorities to complete before 2 August 2026

If a company had to fix three absolute priorities today, these would be mine, in order of urgency.

First, map the AI systems in use – including those hidden inside third-party software – and classify them by role and by Article 50 category. Without this, everything else is blind. The mapping exercise is less technical and more investigative than people expect: start from the business functions rather than from IT, because marketing, communication, HR, customer care and sales are the ones who introduced generative tools autonomously. Then go through procurement and software licences, and ask suppliers explicitly whether their product embeds AI components, whether it generates content, and whether it interacts with users. Record, for each system, the date of placing on the market or putting into service.

Second, implement the disclosures that bite on 2 August: the AI interaction notice in chatbots (Article 50(1)) and the labelling of deep fakes and public-interest texts (Article 50(4)). These are the obligations that arrive without any transitional period, and they are also the most externally visible.

Governance: Deciding who does what

Third, review supplier contracts and set up internal governance. Compliance with the AI Act transparency obligations is not a matter for a single function. Legal and the DPO interpret the obligations and handle the contractual review and the GDPR overlap; IT and the CISO implement the disclosures, the marking and the logs; compliance builds the register and the monitoring; the business owners in each function are responsible for the actual display of the notices; procurement guards the supplier clauses; and senior management provides the mandate, without which governance stays on paper. A small AI governance committee with a responsibility matrix – who does what, for each system – is the model I recommend.

Machine-readable marking under Article 50(2), which can rely on the window to 2 December 2026 for systems already in use, comes immediately after. Important, but with a few extra weeks of breathing room.

Accountability is the real test

The principle I keep repeating to clients is accountability. It is not enough to be compliant; you must be able to prove it, quickly. In concrete terms, I suggest building a compliance file for each system, containing:

ITEM	WHAT IT SHOULD INCLUDE
AI systems register	Role (provider/deployer), applicable Article 50 category, supplier, date of placing on the market or putting into service
Transparency evidence	Screenshots of notices, disclosure wording, label configurations
Technical documentation	Marking and detection solutions, supplier documentation
Logs	Operating logs where required
Internal assessments	FRIA where mandatory, DPIA under the GDPR where relevant, system classification analyses
Contracts	Supplier agreements and evidence of the allocation of obligations

That file is the difference between answering an authority’s request in a day or in a matter of weeks.



AI Act deepfakes and AI-generated text: What deployers must do from 2 August 2026

Author: *Giacomo Lusardi*

On 2 August 2026, the transparency obligations under Article 50 of the EU AI Act become enforceable. The only exception concerns providers of AI systems generating synthetic audio, image, video or text content that were placed on the market before that date. Under the Digital Omnibus Package, those providers have until 2 December 2026 to comply.

Most commentary focuses on the obligations of AI providers, namely the companies developing the technology or having it developed by others. However, Article 50 also imposes separate obligations on deployers, meaning businesses and organizations that use AI systems to generate content.

A quick reminder on providers

Article 50(1) and (2) impose obligations on providers. AI systems designed to interact with individuals must clearly inform users that they are interacting with AI. Providers of generative AI systems must also ensure that synthetic audio, images, videos and text are marked in a machine-readable format.

These obligations, however, do not replace the obligations imposed on deployers.

Deployer obligation #1: Deepfakes

Under Article 50(4), a deployer acting in a professional capacity that uses AI to generate or manipulate an image, audio recording or video must clearly disclose that fact where the content realistically depicts a person, object, place or event and could reasonably be perceived as authentic.

The disclosure must be clear and directly perceivable by humans. Deployers cannot rely solely on the technical watermarking or machine-readable markers embedded by the AI provider.

Examples

Corporate CEO avatar

A company creates an AI-generated video of its CEO delivering a message to employees. The CEO never recorded the message and the video was generated using AI. This is a deepfake and should be clearly labelled as such, even if used only for internal communications.

Advertising campaign

An advertiser uses AI to create a realistic image showing two professional footballers together in a setting where they were never actually present. If the image could be mistaken for a genuine photograph, it must be labelled as AI-generated.

Artistic and satirical content

A lighter transparency requirement applies where the content forms part of an artistic, fictional or satirical work. For example, a clearly humorous AI-manipulated image used for parody would still require a disclosure, but the disclosure should not interfere with the presentation or enjoyment of the work.

Deployer obligation #2: AI-generated text

Deployers must also disclose when they publish AI-generated or AI-manipulated text intended to inform the public on matters of public interest.

Examples may include:

- an AI-generated summary of a sustainability report published on a corporate website;
- an AI-generated summary of a public authority decision;
- AI-generated news or informational content.

The editorial control exemption

The disclosure obligation does not apply where:

- a human has meaningfully reviewed and edited the content; and
- a person or organisation retains editorial responsibility for publication.

A simple grammar or spell check is unlikely to be sufficient. Human review should include substantive oversight and fact-checking.

For example, an article initially drafted by AI but subsequently reviewed, edited and approved by an editor acting under the publisher's editorial responsibility would generally not require an AI disclosure.

Who is in scope?

These obligations apply only to deployers acting in a professional or business capacity.

A private individual creating an AI-generated Christmas greeting for friends would not normally be caught by these requirements, although the provider's obligations under Article 50(2) would still apply to the AI system itself.

The AI Act also contains limited exceptions for certain law enforcement activities.

What should deployers do now?

Before 2 August 2026, organisations should:

- identify where AI systems are used to generate or manipulate images, audio, video or public-facing text;
- adopt internal policies on the use of AI-generated content and on when disclosures are required;
- implement clear disclosure mechanisms and publication workflows; and
- remember that deployer disclosure obligations are separate from, and in addition to, any watermarking or technical measures implemented by AI providers.

Importantly, content generated before 2 August 2026 but published on or after that date may still need to comply with the applicable transparency requirements.

AI Act investigations: How to handle the first information requests

Author: *Giulio Coraggio*

AI Act investigations have begun, with the European Commission using its enforcement powers for the first time to order leading AI developers to detail their practices on cybersecurity, safety and copyright through information requests.

The issue is not only for developers as investigations on whichever company uses artificial intelligence systems will start soon. The way a company answers that first letter is about to become the decisive moment of its AI compliance, because the reply is not a documentation exercise but the opening move of a defence.

Enforcement has stopped being theoretical, since the AI Act provisions that make it possible became operational on 2 August 2026.

Why AI Act investigations reach companies that only use AI

The most persistent misunderstanding I encounter is that the AI Act regulates developers and that everybody else can wait for the deferred deadlines to arrive.

The AI Act does not regulate developers alone, because it places its own set of obligations on deployers, meaning any company using AI systems in its operations, and those obligations have not been postponed. What was deferred concerns parts of the high risk regime and not the framework as a whole, and certainly not the enforcement powers that started to operate at the beginning of August 2026. The transparency duties reach a very wide population of organisations, and the position of a company that customises, retrains or rebrands a system is even less comfortable, since that activity can requalify a deployer as a provider with the heavier set of duties that follows. Treating the deployer role as a shield does not work because the duty follows the system all the way to the user.

The reply is already part of the exposure

There is a feature of the AI Act enforcement framework that most organisations have not internalised, and it changes the nature of the exercise entirely, because liability does not begin with the substance of what a company did with the system. Refusing to respond to an information request, answering it badly or obstructing a review is sanctionable in itself, which means that a letter which looks like a request for documents is in fact the opening of a file, and everything produced in response becomes part of it.

This is why answering an information request is not a question of merely providing technical specifications to be handled by the IT team. The reply has to be drafted with the reasons behind the request in mind and within a defensive strategy agreed from the outset, because the authority is not collecting documents for their own sake, it is testing a potential challenge, and a response that ignores the potential challenge tends to confirm it.

In my experience this is the moment where the final outcome of a dispute is decided, since everything that follows is built on what you said the first time, and no later clarification ever fully repairs an answer given too quickly.

What privacy investigations already taught us

None of this is unprecedented, because data protection authorities have been running this playbook for years, and companies that have been through a privacy inspection recognise the mechanics immediately. I wrote about how to prepare for privacy dawn raids under the GDPR some time ago, and three of the lessons from that experience transfer directly to AI Act investigations.

The first concerns accountability, since under a regime built on that principle the investigated party has to prove that it did what it was meant to do, which broadens the scope of any investigation well beyond the specific question asked and turns the absence of documentation into evidence in itself rather than into a neutral fact.

The second concerns privilege, because in jurisdictions such as Italy legal privilege is not a substantial limit to investigations, and internal exchanges between departments, or in some circumstances even with external counsel, can be used to challenge a company's position, which means that how an organisation writes about its own AI systems internally matters long before anybody asks to see it.

The third, and the one that causes most of the damage in practice, concerns human error, since the recurring source of problems in these files is employees who provide inaccurate information for no better reason than the wish to be helpful and to feel important for a day, and an answer given informally by somebody who does not have the full picture creates a problem that no subsequent legal review can undo.

The procedure that should exist before the letter arrives

The preparation that works is the kind that cannot be improvised once a deadline is running, and it looks very much like the internal procedure that mature organisations already have for data protection inspections, adapted to a different regulator and a different subject matter.

It starts with knowing who the primary contacts are, not only the person responsible for AI governance but also external counsel and a named contact in each office, together with instructions for the people who would physically receive a communication, and with an escalation path that ensures the general counsel and management are informed immediately rather than after somebody has already replied.

It continues with knowing who is entitled to answer on behalf of the company, because an information request that reaches a technical team and gets answered by simply supplying technical details is the single most common way these files go wrong, and the people best placed to explain how a system works are rarely the people best placed to decide what the company should say about it.

It compels knowing which AI systems are actually in operation, including the ones embedded in third party human resources, customer relationship, security and procurement software that nobody classified as artificial intelligence when it was purchased, and including the generative tools that individual teams adopted on their own initiative, because an inventory assembled after a request arrives is both incomplete and inconsistent with whatever the company says next.

It requires knowing what you would be able to show today, since the value of documentation is decided by whether it exists at the moment of the request and not by whether it could be reconstructed afterwards.

And it requires training, which is the part that gets skipped, because a procedure that nobody has rehearsed is a document rather than a defence, and the people who will be interviewed need to understand that they should answer questions whose answer they actually know, refer to the internal policies and procedures where they exist, and refrain from offering personal opinions on matters that will later be read as the company's position.

What comes next

The developers that received the first requests will handle them with counsel in the room from the first hour, and their outcome will turn on questions of substance they have been preparing for years.

The organisations that should be paying attention are the ones that concluded the AI Act was somebody else's problem, because the enforcement model that begins with a letter is the model every digital regulator uses, and it always starts long before anybody talks about a fine. Waiting was a defensible position while enforcement remained theoretical, and it stopped being one when the first letters went out.

Understanding where an organisation stands usually begins with mapping the AI systems in use, the role the company holds in relation to each of them and the documentation that exists around them, which is a defined exercise and considerably less expensive than assembling the same information inside the deadline of a request.

To ease the understanding of the issue, below are the answers to some FAQs:

What are AI Act investigations?

They are the enforcement proceedings through which the European Commission, the AI Office and the national authorities verify compliance with the AI Act, and they typically begin with information requests asking operators to detail their practices, produce documentation and explain how their systems work.

Can a company be sanctioned for how it answers an information request?

Yes, because refusing to respond, answering badly or obstructing a review is sanctionable in itself, independently of whether the underlying conduct concerning the AI system turns out to breach the Regulation.

Do AI Act investigations concern companies that only use AI systems?

Yes. The AI Act places obligations on deployers, meaning the organisations that use AI systems, and those obligations were not postponed, while a company that customises, retrains or rebrands a system may also be requalified as a provider with the heavier duties that follow.

What was actually postponed?

The postponements concern parts of the high risk regime rather than the general application of the Regulation, which took effect on 2 August 2026 together with the enforcement powers of the Commission and of the national authorities.

Is legal privilege a protection during these investigations?

Not to the extent that many companies assume, since in jurisdictions such as Italy privilege is not a substantial limit to investigations, and internal communications about the company's own systems can end up being used to challenge its position.

Who should answer an information request within a company?

The reply should be coordinated by whoever is responsible for the company's legal position, with technical input rather than technical drafting, because the response is evidence and it has to be consistent with the defensive strategy adopted from the outset.

How should a company prepare for AI Act investigations?

By mapping the AI systems in operation including those embedded in third party software, identifying the role the company holds for each of them, establishing the internal contacts and the escalation path, verifying that the documentation that would be requested already exists, and training the people who would be involved.



AI criminal liability in Italy: What boards must know now

Author: *Giulio Coraggio*

AI criminal liability is no longer a question of how much a fine would cost, it is a question of who signs off on a system and what happens to that person, and to the company, if the system causes harm.

Italy has now answered that question, and the answer travels beyond its borders. On 4 August 2026 the Council of Ministers gave final approval to the two legislative decrees adapting Italian law to the AI Act, completing the delegation contained in Law No. 132/2025. What deserves separate treatment is the part that reaches individuals and companies directly, because it changes both how a deployment decision should be taken and who carries the consequences of taking it.

Two things happened at the same moment. A new offence entered the Criminal Code, Article 437-bis, punishing the omission of safety measures in high-risk AI systems, and that offence simultaneously became a trigger for corporate criminal liability, entering the catalogue of predicate offences under Legislative Decree 231/2001 through a dedicated new article devoted to crimes committed with the use of AI systems.

If your organisation operates in Italy, sells into Italy or has an Italian subsidiary, this concerns you, because the exposure follows the conduct and the place where the offence is committed rather than the nationality of the parent company, and Italian corporate criminal proceedings routinely reach entities whose decision makers sit somewhere else entirely.

Why AI criminal liability changes the decision

Administrative fines under the AI Act are serious and they are budgeted for, because a fine is ultimately a cost that the company absorbs, whereas AI criminal liability is faced by a person and no budget line covers it.

That difference reaches the approval stage of every AI project. When the downside was monetary, governance could be weighed against expected benefits and it usually lost the comparison, while now that the downside includes personal exposure for the people who approved the deployment, the question is no longer how much this could cost but whether the person who authorised the system can show what they did before authorising it.

In the incidents I handle, that evidence either exists or it does not, and it is never created afterwards.

What Article 437-bis punishes

The offence covers whoever omits the technical safety measures required to prevent malfunctions or alterations, or omits human oversight measures, across the design, training, production, placing on the market or professional use of high-risk AI systems.

Liability is not triggered by the omission alone, because the provision requires what Italian criminal law calls concrete danger, meaning a real and demonstrable risk to a protected interest in the specific circumstances rather than an abstract possibility that harm might occur.

The penalties then follow the interest at stake, running from one to five years of imprisonment where life or individual safety is exposed and from two to eight years where the danger reaches public safety or the security of the State, while a second limb of the provision deals with outside interference, punishing whoever, remaining external to the system lifecycle, alters its functioning, with two to six years in the base case and three to ten years in the aggravated ones.

Where the omissive conduct is committed with gross negligence the penalty is reduced by between one third and one sixth, on the model of Article 590-sexies of the Criminal Code, which governs negligent liability for death or personal injury in healthcare, the intention being to confine criminal relevance to macroscopic departures from diligence in a field that is changing very quickly and to leave ordinary technical error outside the perimeter.

The placement in the Code says more than the text does, because Article 437-bis sits next to Article 437, which for decades has policed the failure to install equipment designed to prevent workplace accidents, so Italian law has effectively decided that failing to secure an AI system belongs to the same family as failing to install a safety guard on a machine, and anyone who has lived through a workplace safety investigation will know what that comparison implies about how these files tend to be built.

How corporate criminal liability works in Italy

For readers outside Italy this is where the analysis stops resembling most other European systems, so the mechanics are worth setting out.

Legislative Decree 231/2001 established a regime under which a company is sanctioned directly when a specified offence is committed by its senior managers or its employees in the interest or to the advantage of the company, and only offences included in a defined statutory list, known as predicate offences, can trigger it, which is why adding an offence to that list is a significant legislative act and why what has just happened to AI matters.

Two features deserve attention from anyone used to a different system. The sanction is imposed on the entity itself, in its own proceedings, alongside any case brought against the individuals involved, and the company has a genuine defence rather than mere mitigation where it can show that it had adopted and effectively implemented an organisational, management and control model designed to prevent offences of that kind, overseen by an independent supervisory body. The practical implication is direct, because a model that does not address AI risk cannot function as a defence against an AI predicate offence.

On sanctions, the new provision applies a monetary penalty of between six hundred and one thousand units for Article 437-bis and between two hundred and seven hundred units for the deepfake offence under Article 612-quater introduced by Law 132/2025, and since Italian corporate sanctions are expressed in units rather than in euro, with the court setting the monetary value of each unit within a statutory range according to the financial condition of the entity, the same number of units produces very different amounts for a small supplier and for a multinational, reaching well above one million euro at the upper end.

Where the legislator positioned the offence tells you more than the numbers do, because the explanatory report accompanying the decree places it in the high band of the catalogue and adopts as its express benchmark the predicate offence covering manslaughter and personal injury caused by breaches of workplace safety rules, which means that AI safety failures have been ranked, quite deliberately, alongside industrial safety failures.

Interdictive sanctions apply as well and in practice they hurt considerably more than the money, since they include the suspension or revocation of authorisations, the prohibition on contracting with the public administration, the exclusion from grants and financing and the prohibition on advertising goods or services, while interdiction from carrying on the business activity was left out on proportionality grounds, which is the one piece of good news in the package.

Why deployers are inside the perimeter, not outside it

This is the point that most companies get wrong when they read the headline.

The conduct covered by Article 437-bis runs across design, training, production, placing on the market and professional use, and that list is not accidental because it tracks the value

chain of the EU AI Act from provider through to deployer and distributor, with the consequence that AI criminal liability, and the corporate exposure that follows from it, reaches every company that professionally uses high-risk AI systems developed by somebody else.

Most organisations I speak to assume that criminal exposure for AI belongs to whoever built the model, which is not the case, because if you deploy a high-risk system in your operations the duty to maintain safety measures and human oversight is yours and remains yours for as long as the system runs, and buying that system from a reputable vendor moves nothing at all.

One consequence deserves attention on the procurement side, since if your exposure depends on measures that only the provider can implement and on information that only the provider holds, then your contracts need to secure both, and most AI supply agreements signed in the last two years do neither.

What changes when the system acts on its own

A recommendation engine proposes, whereas an agent books, negotiates, executes and decides, increasingly without a human approving each step.

Read Article 437-bis with that shift in mind and the exposure becomes clearer, because the offence is built around two elements that agents complicate, the first being the adequacy of the technical measures preventing malfunction and the second being human oversight, which the provision treats as an autonomous duty whose omission is punishable in itself, so that an agent operating with wide autonomy and no meaningful oversight point is, in the terms of this provision, the paradigm case.

The investigative consequence follows naturally, because the chain between a management decision and a harmful outcome becomes longer and more autonomous, and the question of who authorised what stops being an organisational detail and becomes the centre of the case. If an agent takes an action creating concrete danger, the first requests will concern who decided to deploy it, what testing was performed on its impact, what limits constrained what it could do, and who was monitoring it when the situation deteriorated.

I advise a growing number of organisations that are introducing agents into their operations and in many of them none of those four questions has a documented answer, because systems reach production without prior impact testing, without guardrails and without any governance model capable of monitoring how they behave over time, which under a regime of administrative fines was a commercial risk and under a regime of AI criminal liability is a different order of decision, usually one that nobody in the organisation has consciously taken.

The interest or advantage test, and where it points

Corporate liability requires that the offence was committed in the interest or to the advantage of the company, and that connecting factor, which prosecutors have to establish, is what links an individual's conduct to the entity.

On this particular offence the advantage is easy to identify and that is precisely the problem, because it consists of the saving on safety measures, meaning the controls not implemented, the testing not performed, the oversight function not staffed and the monitoring never built.

The structure is uncomfortable, because the same organisational failure risks being counted twice, once in the conduct itself and once again in the organisational fault that founds corporate liability, and Italian commentators have flagged the point alongside a textual overlap in the first paragraph where public safety appears in both the base and the aggravated bands, neither of which was resolved at the parliamentary stage and both of which will be argued.

For a board the practical reading is simpler, since any decision to defer AI safety spending now produces a documentary trail pointing directly at the advantage element of a corporate criminal case.

One complication worth holding in mind

The offence attaches to high-risk AI systems and the European framework for high-risk systems has just moved, because the Digital Omnibus, Regulation (EU) 2026/1744 of 8 July 2026 published on 24 July, postpones the application of the high-risk provisions while leaving the rest of the AI Act in place.

The tempting conclusion is that there is time, but the safer reading is that the classification exercise cannot wait, because a company that has not established which of its systems are high-risk cannot know which of them sit inside Article 437-bis and will not be able to show, later on, that it ever looked.

What to put in place

None of the protective measures are exotic and most of them make the AI project work better in any event.

Map the AI systems in use and classify them, since inventories are usually the missing piece, particularly for AI embedded in third party HR, CRM, security and procurement software and for generative tools that staff adopted on their own initiative.

Test the impact before the system goes live and keep the assessment, covering what the system can do, what can go wrong and who could be harmed, because the value of that document lies not in the analysis itself but in the proof that the analysis happened.

Build the human oversight function and give it both a name and the authority to stop the system, since the provision treats missing oversight as punishable conduct in its own right, which makes this the single measure with the clearest link to the offence.

Set the boundary between autonomy and authorisation in writing, defining what an agent may do alone and what requires a human decision, because systems inherit the limits that somebody defined for them and where nobody defined any, the limits are whatever the model happens to produce.

Fix the supply chain contracts on security measures, oversight support, incident information, audit rights and cooperation in a proceeding, because your exposure depends on things that your provider controls.

Update the compliance model that operates as a defence, which for groups with Italian operations means revisiting the local organisational model rather than the global policy suite, since the defence depends on the Italian document and on evidence that it was actually implemented and supervised.

Organisations that want to understand where they stand on AI criminal liability usually start with an assessment of the AI systems in use and of the governance around them, which is a defined exercise with a defined scope and costs a fraction of what it costs to reconstruct the same information under pressure.



CADA: The European strategy for cloud and AI between industrial support, sovereignty and public procurement.

Author: *Giulio Napolitano*

On 3 June 2026 the European Commission presented the proposal for a regulation COM (2026) 502 final establishing a framework of measures for the strengthening of the European *cloud* and artificial intelligence ecosystem (hereinafter, the *Cloud and AI Development Act* or “CADA”). The context is well known, but it is worth recalling. The market share of European providers of *cloud* services fell from 29% in 2017 to 15% in 2022, and then stagnated, while three non-EU *hyperscalers* today control over 70% of the European market. On top of this dependence there arises a risk that is not only economic but also legal, since the dominant operators are subject to third-country legal systems whose rules, having extraterritorial effect, may impose data access and transfers conflicting with the fundamental rights of the Union.

The systemic novelty of CADA lies in where it places sovereignty. With the GDPR, European autonomy had been built on personal data. With CADA it descends one level and takes root in the infrastructure, that is, in the ownership, corporate control and physical location of servers, personnel and the software chain. It is not a *cybersecurity* intervention, which remains entrusted to other instruments, but rather the codification of a structural criterion of autonomy. The regulation rests on a dual legal basis, Article 114 TFEU on the internal market and Article 173(3) TFEU on industrial policy, and pursues two declared objectives, the competitiveness of the *cloud* and AI ecosystem and the strengthening of the Union's resilience and strategic autonomy.

1. A regulation with a dual soul: support for supply and the lever of public demand

CADA combines supply-side support measures and demand-side stimulus measures.

- On the supply side, Title II establishes the *Cloud and AI Leadership Initiatives*, articulated into eight *grand challenges* (sustainability and energy efficiency of *data centres*, *cloud stack* end-to-end, *frontier AI*, *physical AI*, *industrial AI*, cooperative industrial models, an AI agent platform and public-sector AI). To these are added a mechanism for recognising priority *frontier AI* projects (Articles 8 and 9) and the obligation for Member States to adopt, within one year, a national strategy consistent with the regulation (Article 7).
- Title III addresses the computing capacity gap. Member States designate *data centre acceleration zones* with simplified authorisation procedures and one-stop shops,

while the Commission may recognise *data centre strategic projects* with strong innovation and sustainability content. The declared ambition is to triple the Union's capacity within five to seven years and to cover estimated needs by 2035, with a balanced geographical distribution. This industrial pillar, though significant, is the frame. The prescriptive core of the regulation lies elsewhere, in Title IV.

2. The four assurance levels: sovereignty becomes gradable and certifiable

Title IV establishes a *Union cloud computing sovereignty framework* founded on four *Union assurance levels*, whose cumulative criteria are set out in Annex II. The logic is one of increasing intensity.

1. **Level 1** requires establishment in the Union, the location of infrastructure and data within the territory of the Union (unless the public contracting authority requests otherwise), full transparency on subcontractors and, for providers controlled by a third country, the guarantee that no law of that country requires the disclosure of software vulnerabilities to foreign authorities before their known exploitation.
2. **Level 2** adds the European *cybersecurity* certification of at least a substantial level, the prohibition on using the data generated by the service to train third parties' AI systems, and measures of legal, technical and organisational separation from any non-EU parent companies.
3. **Level 3** raises the bar further and, as a rule, excludes that the provider and the subcontractors involved be subject to the control of a third country, requiring personnel with Union citizenship and, where appropriate, security *clearances*.
4. **Level 4** reserved for sensitive data identified following a risk assessment, requires effective control that is entirely European, personnel composed solely of Union citizens and a high-level *cybersecurity* certification. The decisive criterion is the *test of ownership and control* detailed in Annex III, which requires tracing the ownership chain up to the beneficial owners, mapping the *cap table*, the veto rights, the decision-making quorums and the commercial or financial ties capable of conferring de facto control, with a materiality threshold set at 5% of the capital or of the voting rights.

Recognition is governed by Article 17. The provider submits an application to the competent national authority of establishment, which assesses the evidence within 60 days and triggers a review period of a further 60 days open to the reasoned objections of the other national authorities, with possible referral to the Commission in the event of persistent disagreement.

For Level 1, conformity is attested by means of a *self-assessment* and an EU declaration of conformity (Article 19), with automatic and direct recognition in all Member States for SMEs.

For Levels 2 to 4, an independent third-party *audit* is instead required, with a positive or negative opinion, annual review and strict independence requirements for the *audit* body (Articles 20 and 21). The recognised services flow into a public *central repository* managed by the Commission (Article 22).

3. Associated third countries: an enhanced adequacy for Level 3

Openness towards the outside is governed by Article 18. The Commission may adopt implementing acts identifying the third countries whose controlled providers may nonetheless be subjected to an *audit* for Level 3, provided that cumulative criteria are met.

The third country must be:

- the addressee of an adequacy decision pursuant to Article 45 GDPR,
- it must not have measures allowing control that conflicts with the rules on lawful access to non-personal data under Article 32 of the Data Act,
- it must not be able to impose the interruption of the service nor compliance with sanctions regimes or embargoes (unless these are legitimate under the law of the Member States or of the Union),
- it must not hinder the supply of state-of-the-art technologies, it must keep a market open to the Union's *cloud* services and guarantee equivalent access to its own procurement procedures (reciprocity clause).

This is, in essence, an enhanced adequacy, which adds to the protection of personal data an assessment of infrastructural autonomy and commercial reciprocity.

4. The demand lever: risk assessments, mandatory levels and Union added value

The second pillar of Title IV turns public spending into an instrument of industrial policy. Within one year, and then every two, Member States and Union institutions carry out risk assessments identifying the public-sector activities relevant to public order, namely those in the sectors of Annexes I and II of the NIS2 Directive and in the areas of national and internal security, border management, defence, justice and law enforcement, determining the appropriate *assurance* level (Article 29). The methodology is set by

the Commission by means of implementing acts, and the Commission may impose a different level where it considers the one identified by the Member State to be inadequate.

- Article 30 translates the assessment into an obligation. Administrations whose activities are not relevant to public order must nonetheless procure services recognised at least at Level 1, while those relevant to public order may procure only services recognised at Levels 2, 3 or 4. The derogations are confined to the cases in which no recognised service is available in the *central repository*, in which a similar tender received no bids in the previous year, or in which the application of the regulation would entail disproportionate costs.
- Private entities falling within Annex I of NIS2 may carry out similar assessments on a voluntary basis (Article 31), without prejudice to the Commission's power to make them mandatory for highly critical sectors.
- To this is added the criterion of *Union added value* (Article 32). In procedures for innovative *cloud* services and AI systems, administrations include non-price award criteria that value the tenderer's contribution to the development of the European ecosystem, such as the use of *software* or *hardware* designed or manufactured in the Union, the integration of European technologies and research results, and the strengthening of security of supply. The regulation expressly requires that such criteria remain ancillary and not decisive for the purposes of the award.

Completing the picture are the objective of allocating at least 25% of *cloud* and AI procurement to innovative SMEs (Article 33), the establishment of a public-sector *cloud* federation (*EuroCloud Federation*, Articles 34-36) and a marked push towards *open source* solutions (Articles 41-44).

5. Governance and penalties

Enforcement is entrusted to competent national authorities, which the Member States designate within one year and which enjoy exclusive competence in the State of the provider's main establishment (Article 25).

- They are vested with significant investigative and *enforcement* powers, from access to premises to the order to cease the infringement, to penalties and periodic penalty payments (Article 26).
- Penalties are defined by the individual Member States, according to the usual standard of effectiveness, proportionality and dissuasiveness, and among the criteria for setting them is the offender's annual turnover in the Union.

Service recipients are further granted a right to compensation for damage arising from breaches of the provider's obligations (Article 24). Consistency of application is safeguarded by mutual assistance mechanisms and by a Commission power to adopt binding decisions in the event of disagreement between authorities.

6. The critical knots: the internal market against itself

The construction raises significant systemic questions.

1. The first is dogmatic. A regulation founded also on Article 114 TFEU, whose function is to remove obstacles to the internal market, erects a regime that discriminates among providers on the basis of ownership and corporate control, that is, precisely one of those criteria that competition law and the freedom to provide services tend to neutralise. It is no coincidence that the regulation takes care to specify that it does not prejudice the application of Articles 101 and 102 TFEU. There remains, however, the tension with technological neutrality and with the free movement of services, which constitute the genetic code of the single market.
2. The second question is international. The reciprocity clause of Article 18 and the *Union added value* criteria of Article 32 move at the margins of the obligations assumed by the Union in the field of public procurement, in particular under the WTO *Government Procurement Agreement*. It is not surprising that, in the aftermath of the presentation, trade associations from third countries expressed the fear of a de facto exclusion from the most critical tenders. The characterisation of the *Union added value* criterion as ancillary and not decisive is the valve that should preserve the compatibility of the system, but its resilience will be measured in the practice of tenders, where even a modest differential can reshape the outcomes.
3. The third question concerns the coherence of the digital *acquis*. CADA is added to an already dense regulatory fabric (e.g. GDPR, Data Act, NIS2, DORA, AI Act, revision of the Cybersecurity Act) precisely while, with the *Digital*

Omnibus, the Commission declares its intention to reduce regulatory layering. The addition of a further level of sovereignty, however oriented towards strategic autonomy, will require careful coordination with the *cybersecurity* certifications and with the obligations on third-party ICT providers, in order to avoid duplications and application frictions.


7. What to do now

On an operational level, public contracting authorities will have to map their own activities to verify their relevance to public order, identify the required *assurance* level and plan any migrations, which must be completed within a period that, as a rule, does not exceed twelve months.

The *cloud* providers will have to choose the level to aspire to, distinguishing between *self-assessment* for Level 1 and an independent *audit* for the higher levels, and to begin in good time the reconstruction of their own *ownership and control* chain.

Private entities subject to NIS2 will have to assess the appropriateness of carrying out the voluntary analyses of Article 31.

A final note concerns the *fintech* sector and providers of *crypto-asset services*. Many regulated operators rely on the infrastructure of non-EU *hyperscalers*, and CADA intercepts this dependence from the sovereignty side where DORA already governs it from the side of the ICT risk of third-party providers. The coordination between the two regimes will deserve specific attention, in the context of contractual *due diligence* and of *multi-cloud* strategies. Since this is a text at an early stage, every numerical reference to articles and annexes will need to be verified against the consolidated text, which may change in the course of the legislative process.



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