SCRIPTORIUM INITIATIVE

USING AI WITHOUT LOSING TRUST

A practical guide to the EU AI Act for communications and brand leaders | August 2026

MAKING SENSE OF THE *[_UNKNOWN]*

SCRIPTORIUM-INITIATIVE.AI

EXECUTIVE SUMMARY

The EU Artificial Intelligence Act is a European law governing how artificial intelligence is developed, offered and used. Its purpose is to support innovation while protecting safety, fundamental rights and public trust. The central principle is simple: the greater the potential harm, the stricter the rules.

The Act prohibits specified uses that the EU considers unacceptable. It places demanding safeguards on legally defined high-risk systems, including some AI used in employment, education, essential services and regulated products. It also introduces transparency duties for situations in which people should know that AI is involved. Most ordinary, low-impact uses face few or no special obligations under the Act, although other laws still apply.

For communications and brand leaders, Article 50 is especially important. It addresses direct interaction with AI, technical marking of AI-generated content, certain biometric and emotion-recognition systems, deepfakes and some AI-generated text published to inform the public on matters of public interest.

> **GLOBAL REACH** A company does not need an EU headquarters to be covered. A US or other non-EU company may fall within the Act when it offers AI in the EU or when output produced by its AI system is used in the EU.

What leaders should remember

- Know where AI appears across customer experiences, content workflows, agencies and platforms.
- Understand what the AI is doing; the use case determines the regulatory risk.
- Make required disclosures clear, immediate and accessible.
- Keep a person accountable for review, approval and correction.

WHY THE ACT MATTERS TO GLOBAL ORGANISATIONS

AI is no longer only a technology issue. It influences decisions, customer interactions, corporate information and brand trust. The AI Act creates a common European framework for determining when AI should not be used, when strong safeguards are necessary and when people should be told that AI is part of an experience or item of content.

Communications and brand teams sit where these expectations become visible. They decide what an assistant calls itself, how campaign assets are labelled, how public statements are reviewed and how incidents are explained. A technically compliant system can still damage trust if audiences experience it as hidden, misleading or unaccountable.

Does the Act apply to us?

Start with where the AI is offered and where its output is used—not where the company is headquartered. The Act can cover organisations outside the EU when they place an AI system or general-purpose AI model on the EU market. It can also cover providers and deployers outside the EU when output produced by the system is used in the EU.

An EU office using an AI system may create responsibilities for the organisation. Working through a vendor or agency does not automatically remove them. Each party's role and the facts of the use must be understood.

> **EXAMPLE** A US retailer offers its branded AI shopping assistant to customers in France and Germany. Its US headquarters does not keep the activity outside the Act. The precise responsibilities depend on how the assistant is developed, branded and supplied.

Two roles in everyday language

A provider generally develops an AI system—or has it developed—and offers it under its own name. A deployer uses an AI system under its authority. The same company can be a provider in one project and a deployer in another.

UNDERSTANDING THE RISK FRAMEWORK

The Act is often explained as a risk pyramid. The useful idea behind that image is that the rules follow the purpose and likely effect of the AI, not simply the power of the tool. The same model might help draft a routine email in one setting and support a sensitive employment decision in another.

Banned uses

Article 5 prohibits specified practices considered incompatible with EU values and fundamental rights, including certain forms of harmful manipulation, exploitation and social scoring. A warning label cannot make a prohibited practice acceptable. The response is to stop and escalate.

High-risk uses

Some AI can significantly affect safety, rights or access to important opportunities. Legally defined high-risk examples include certain systems used in recruitment, education, credit, essential services, biometrics and regulated products. Requirements can cover risk management, data governance, documentation, logs, human oversight, accuracy, robustness and cybersecurity.

Transparency-focused uses

Some systems are regulated because people need to understand that AI is involved. Chatbots, realistic synthetic media and certain AI-generated public-information text are especially relevant to communications teams.

Everyday uses

Most ordinary AI uses face few or no special requirements under the AI Act. Spam filters, many productivity tools and low-impact content assistance may fall here. Privacy, consumer, copyright, employment and other laws can still apply, as can reputation risk.

> **DO NOT CONFLATE THE LAYERS** A chatbot may be low-risk in one context and still require an AI notice. Transparency obligations can apply alongside other duties.

ARTICLE 50: THE TRANSPARENCY RULE

Article 50 is the part of the Act most likely to appear directly in communications and brand work. It does not require every use of AI to be labelled. Instead, it identifies situations in which transparency protects people from confusion or deception.

When people interact directly with AI

People should be informed that they are interacting with AI at the beginning, unless this is already genuinely obvious in the circumstances. A friendly name alone may not make an assistant's artificial nature clear.

When AI systems create content

Providers of systems that generate audio, images, video or text may need to place a machine-readable signal in the output. This technical marker allows systems to detect that content was artificially generated or manipulated.

Sensitive recognition systems

People exposed to emotion-recognition or biometric-categorisation systems must generally be informed that the system is operating. Data-protection requirements apply alongside the AI Act.

Deepfakes and public-interest text

Realistic synthetic images, audio or video generally require disclosure. AI-generated or manipulated text published to inform the public on matters of public interest may also require disclosure unless meaningful human review or editorial control exists and a person or organisation holds editorial responsibility.

TWO FORMS OF TRANSPARENCY Technical marking is designed for machines; a visible disclosure is designed for people. One does not automatically replace the other.

BOTS AND AI ASSISTANTS

A chatbot is often the clearest example of Article 50 in practice. The central question is whether a reasonable person understands from the beginning that they are interacting with AI. A branded name such as “Ask Ava” may sound like a human service agent; an “AI shopping assistant” label removes doubt.

What good looks like

- Identify the assistant as AI in the opening message and keep the identity visible where practical.
- Use plain language rather than vague terms such as “digital colleague.”
- Explain important limitations, particularly where an incorrect answer could matter.
- Provide a route to a person when service expectations, policy or risk require it.
- Test the notice for accessibility and retain the approved wording.

SAMPLE OPENING “You’re chatting with Ava, our AI shopping assistant. It can help you explore products, but its answers may be inaccurate. For account or complaint support, connect with our team.”

Match the wording to the context

A low-stakes product finder needs less explanation than an assistant discussing financial, health or employment matters. The legal duty is to identify the interaction with AI unless it is already obvious. A company’s brand standard can go further when trust or potential harm warrants it.

Before launch

Product and communications should confirm the opening notice. Brand and accessibility teams should test the language. Operations and legal should verify the stated limitations and human-support route. The product owner should ensure the notice survives updates and localisation.

AI-GENERATED IMAGES, VIDEO AND VOICE

The Act defines a deepfake as AI-generated or manipulated image, audio or video that resembles real people, objects, places or events and could wrongly appear authentic. When an organisation publishes this kind of content, it must generally disclose that the material was artificially generated or altered.

The more realistic the content, the greater the care

A clearly fantastical campaign visual is unlikely to mislead in the same way as a realistic synthetic person. Even so, teams should preserve any technical AI marker and consider whether a voluntary label would help the audience. A realistic person who does not exist needs clearer treatment because viewers may assume that the person and scene are authentic.

An executive's cloned voice or likeness is more sensitive still. Obtain explicit permission, disclose the use prominently and prepare for possible impersonation or misuse. Minor editing of a real photograph may be treated differently when it does not substantially change the image or its meaning. Artistic, satirical and fictional works receive tailored disclosure treatment.

> **SAMPLE DISCLOSURE** "This video contains an AI-generated version of our spokesperson's voice and image."

A reliable publishing pattern

- Place the notice with the content, not only in terms and conditions.
- Make it visible or audible before the content is likely to mislead.
- Keep the label when content is cropped, shared or republished.
- Do not remove provenance information added by the AI tool.
- Record permission, tool, material edits, label and final approval.

AI-ASSISTED ARTICLES AND STATEMENTS

Article 50 addresses AI-generated or manipulated text published to inform the public about matters of public interest. Disclosure is generally required, but an exception may apply when the content receives meaningful human review or editorial control and a person or organisation holds editorial responsibility.

What meaningful review looks like

- A named editor checks claims against reliable primary sources.
- The editor changes or rejects material that is inaccurate, misleading or inappropriate.
- Legal or subject-matter review is obtained where the topic requires it.
- The organisation accepts responsibility for the finished publication.
- The review, changes and approval are recorded.

NOT ENOUGH ON ITS OWN A generic “reviewed by a human” checkbox does not demonstrate meaningful editorial control.

Example: an AI-assisted policy explainer

A company asks AI to produce a first draft explaining its climate-policy position. A named editor verifies every claim, rewrites key sections, obtains expert and legal review and accepts editorial responsibility on behalf of the company. Legal counsel concludes that the editorial-control exception applies. The company may still disclose AI assistance if readers would reasonably consider it important.

VOLUNTARY WORDING “AI was used to support the first draft. The final article was fact-checked, edited and approved by our corporate affairs team, which is responsible for the content.”

WHAT COMMUNICATIONS LEADERS SHOULD DO NOW

A workable governance model does not need to begin with a large new programme. It begins by making AI visible to the organisation and attaching clear decisions and owners to the most important uses.

Five actions

- Find it: list audience-facing AI, including tools operated by agencies and vendors.
- Classify it: decide whether the use is banned, high-risk, transparency-focused or ordinary.
- Disclose it: use approved wording, placement and accessibility rules.
- Review it: name the person responsible for checking important content and approving publication.
- Record it: keep the output, sources, important edits, disclosure, provenance and approval.

The leadership test

Can the organisation explain where AI was used, what people were told, who checked the result and who accepted responsibility? If the answer is unclear, the process is not ready.

Primary references

Regulation (EU) 2024/1689 (Artificial Intelligence Act): eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689

Regulation (EU) 2026/1744 amending the AI Act: eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32026R1744

European Commission AI Act overview: digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

AI Act Service Desk, Article 50: ai-act-service-desk.ec.europa.eu/en/ai-act/article-50

This paper is an accessible executive overview, not legal advice. Confirm the current consolidated text and obtain counsel for a specific system or market.

MAKING SENSE
OF THE *[_UNKNOWN]*

SCRIPTORIUM-INITIATIVE.AI