



ACF IN ACTION

acfstandard.com

ACF en Action applique l'Agentic Commerce Framework® à des entreprises réelles, à partir de sources publiques. Chaque cas montre qui reste responsable quand l'IA décide.

ACF-CS-023 · Case Study #023 · Édition 2026

EU AI Act (article 50)

How do you guarantee that a person always knows they are interacting with an AI, or that a piece of content was generated by an AI?

CARTE D'IDENTITÉ

SECTOR Regulation	COUNTRY European Union	TECHNOLOGY AI transparency	TYPE Framed	AUTONOMY framed
CONFIDENCE High				

OUTILS ACF ILLUSTRÉS DANS CE CAS

● ACF-01	● ACF-03	● ACF-05	● ACF-08	● ACF-00	● ACF-02
----------	----------	----------	----------	----------	----------

● mobilisé · ● à prévoir · 6 outils du Toolkit (4 mobilisés, 2 à prévoir, sur 17).

1. Executive summary

The European Union adopts the AI Act, the first general framework for the regulation of artificial intelligence. Its article 50 sets transparency obligations for AI systems that interact with the public or produce content. The delegated decision is that of any conversational agent or content generator that addresses people. The approach is distinctive: instead of leaving the user unaware of the nature of their interlocutor, the law requires it to be revealed.

Read through the lens of ACF, this case is a reference case on the transparency of an agent. Where an opaque agent speaks with authority without revealing that it is a machine, article 50 requires the user to always know what they are dealing with. The value of ACF here is confirming: the law materializes, in law, what the framework calls a constitutional rule on transparency. Trust does not rest on "the AI does not deceive", but on the fact that the user knows what they are dealing with and can adjust their judgment.

2. Context

As conversational agents and content generators spread to the public, a governance question becomes central: does the user always know they are speaking to a machine, or that a piece of content was produced by an AI? Without a marker, they may receive a generated response as the word of a human expert, or take an artificial piece of content for authentic. The AI Act answers with a constraint: transparency is no longer left to the publisher's goodwill, it becomes an obligation.

3. Public sources used

- EU Artificial Intelligence Act, "Article 50: Transparency Obligations" (artificialintelligenceact.eu).
- European Commission, "Transparency obligations under Article 50 of the AI Act".
- Legal analyses (Stibbe) on the timeline and scope of article 50.

4. Reconstructed AI architecture (from public sources)

Mechanism imposed by article 50 of the AI Act (obligations applicable as of August 2, 2026).

- Nature: a legal transparency framework for AI systems that interact with the public or produce content.
- Decision concerned: any conversational agent or content generator addressing people.
- Control function: to make the AI nature of the interlocutor or the content visible, rather than to prohibit it or leave it to be guessed.

The documented obligations:

- a conversational agent must flag itself as a machine: the user must know they are speaking to an AI, except where it is obvious;
- AI-generated content (text, audio, image, video) must be marked in a machine-readable format and detectable as artificial;
- deepfakes must be disclosed as generated or manipulated;
- when AI-generated text is published to inform the public on matters of public interest, its artificial origin must be disclosed.

What the mechanism makes possible:

- to know whom you are speaking to: the user knows it is an AI;
- to gauge information: they adjust their trust and verify;
- not to be deceived: the artificial nature is not hidden.

5. ACF mapping

 **Tools used:** ACF-03 (agentic constitution) · ACF-05 (disclosure supervision) · ACF-01 (decision map).

In ACF language, the mechanism reads as follows:

ACF element	In article 50 of the AI Act
Named human (accountable)	The provider or the deployer of the AI system: it is they who answer for the transparency of the agent or the content.
Agent	Any conversational agent or content generator that addresses people.
Platform	The underlying AI system that produces the interaction or the content.
Delegated decision	To interact with the public or produce content intended for people.
Autonomy level	Bounded: the agent may interact and generate, but must reveal its AI nature and mark what it produces.
Non-delegable zones	Leaving the user unaware that they are speaking to an AI, or distributing generated content without marking it: the law prohibits it.
Rule / safeguard	Flagging of the AI nature, marking of generated content, disclosure of deepfakes and of texts informing the public.
Control third party	The national supervisory authorities and the European framework for implementing the AI Act.
Traceability	The machine-readable marking of generated content: it constitutes the verifiable trace of their artificial origin.

A schema in ACF graphic language accompanies this case (separate file).

6. Analysis

 **Tools used:** ACF-03 (transparency rules) · ACF-05 (disclosure supervision) · ACF-01 (accountability).

What the mechanism does well (revealing rather than leaving to be guessed):

- **Transparency becomes an obligation, not an option.** The AI Act does not ask you to trust the AI: it requires revealing its nature. Left to the publisher's goodwill, transparency disappears under commercial pressure; made mandatory, it protects the user.
- **It bears on the nature, not on the performance.** An excellent AI must still flag itself. Disclosure does not judge the quality of the agent, it informs the user of what they are dealing with, a condition of informed consent.
- **It makes the user the judge of what they receive.** Knowing that you are speaking to an AI changes the reception: the user verifies more, does not take the response for the word of a human expert. Marking a generated piece of content allows it to be judged correctly rather than taken for authentic.

What ACF adds and formalizes:

- **Constitutional rules.** Flag itself as AI at the start of the interaction, mark any generated content, disclose deepfakes: this is the translation, into ACF language, of the legal obligations into permanent and signed rules.
- **A disclosure supervision.** Verify, by sampling or automatically, that each interaction begins with the flagging and that the content carries its marking. An unsupervised rule can break silently once coded.
- **A situated accountability.** The agent produces, but the provider or the deployer answers for the transparency: disclosure is not delegated to the system, it engages the one who places it on the market.

7. ACF toolkit used

This grid indicates the tools that the ACF analysis uses (✓) to read this case, and those to link (→
SOON). It does not mean that a given actor uses ACF: it shows which instruments the framework would bring.

Tool	Status	Role in this case
ACF-03 · Agentic constitution	✓	Translate the obligations into permanent transparency rules
ACF-05 · Supervision	✓	Continuously verify the flagging and the marking of content
ACF-01 · Decision map	✓	Locate the agent, the accountable party, and the interaction decision
ACF-08 · Decisions register	→ SOON	Keep the trace of the disclosures and markings carried out
ACF-02 · Criticality matrix	→ SOON	Calibrate disclosure according to the stake (deepfake, information to the public)
ACF-00 · Sovereignty Score	→ SOON	Assess the maturity of the transparency mechanism

8. Governance questions (what an AI committee should ask itself)

1. Why does knowing you are speaking to an AI change the way you receive its response?
2. Is marking a generated piece of content a hindrance, or a condition of trust?
3. Should transparency be left to the publisher's goodwill, or be an enforceable rule?
4. Do our agents flag themselves as AI at the start of each interaction with the public?
5. Do we continuously verify that the flagging and the marking work, or do we assume them acquired once coded?

9. General lessons

The transparency of an agent is not an extra added if time permits: it is the condition for the user to be able to judge what they receive. The AI Act does not constrain the AI, it requires that its nature be revealed, that its content be marked, that its deepfakes be disclosed. What grounds trust is not the promise that an AI does not deceive, but the fact that the user knows whom they are speaking to and can adjust their judgment.

The transferable lesson: an agent that addresses people must reveal that it is an AI, and generated content must be marked. Made mandatory, transparency protects the user; left optional, it disappears.

10. What ACF would do (reference design)

Prescriptive and non-critical formulation: here is how an organization would design such an agent directly with the framework, whatever its sector.

If an organization wished to deploy a public agent directly with ACF, it could in particular:

- **declare** in a signed agentic constitution (ACF-03) the permanent rules: the agent flags itself as AI at the start of the interaction, any generated content is marked, deepfakes and content informing the public are disclosed;
- **organize** a disclosure supervision (ACF-05): verify, by sampling or automatically, that each interaction begins with the flagging and that the content carries its marking, with an alert in case of failure;
- **map** the actors (ACF-01) to recall that the provider or the deployer, not the agent, answers for the transparency;
- **calibrate** disclosure according to criticality (ACF-02): a deepfake or a text informing the public calls for reinforced disclosure;
- **document** a register (ACF-08) keeping the trace of the disclosures and markings, so as to be able to answer for them.

None of this is a reproach: it is the way ACF would make **explicit, signable, and auditable** what the law already imposes, by embedding transparency in continuous governance rather than in a box checked once.

11. Similar cases (transferability)

This case is specific neither to Europe nor to a single sector. It is transferable to any organization that exposes to the public a conversational agent or a content generator:

-  software publishers and SaaS services
-  e-commerce and retail (support and FAQ)
-  media and content platforms
-  marketing and advertising (generated content)
-  banks, insurance, and telecoms (customer service)
-  public services and administrations

Everywhere, the same pattern: an agent that addresses people, an AI nature to reveal, a generated piece of content to mark. This is what makes this document a reference case: it illustrates a concept, not just an article of law.

12. Confidence of the assessment

High. The case rests on an adopted law (EU AI Act, article 50) and its official documentation (European Commission, reference portal), complemented by legal analyses. The obligations and the timeline are therefore established and verifiable by the reader. The ACF analysis draws on these public facts; it bears on the governance of transparency, not on the technical modalities of implementation, still being specified and with no bearing on the reading proposed here.

Teacher guide

Pedagogical objective. Convey that making transparency mandatory transforms trust: the user knows what they are dealing with, independently of the quality of the agent.

Session outline (90 min).

1. (15 min) Reading of the **two** cases in the kit: Cursor (incident) and AI Act art. 50 (reference), without the analyses.
2. (20 min) Workshop: translate the obligations of article 50 into agent rules (card ACF-03).
3. (20 min) In groups: how to verify disclosure continuously (card ACF-05)?
4. (20 min) Comparison Cursor / AI Act: opaque agent that speaks with authority vs the obligation to flag itself.
5. (15 min) Synthesis by the teacher: transparency is an obligation, not an extra; the user must know whom they are speaking to.

Possible exam questions. See the "Governance questions" section.

Reminder: grading students' work is a pedagogical act of the teacher. This guide provides the reference analysis, not automatic grading.

Associated lab (ACF-LAB-023)

From this case, the student produces:

- the constitutional rules (card ACF-03): flag itself as AI, mark generated content, disclose deepfakes;
 - the disclosure supervision mechanism (card ACF-05): verification by sampling or automatic, alert in case of failure;
 - the mapping of actors (card ACF-01): agent, accountable provider or deployer, interaction decision;
 - a one-page comparison between an opaque agent (Cursor) and an agent subject to the transparency obligation.
-

Appendix - Completed ACF cards

The official Toolkit cards used in this case, filled in with its public data. For illustration: an organization of the same profile can reuse and adapt them.



OBJECTIVE

Map the decisions of a public-facing agent subject to Article 50 and set, for each, the autonomy level based on what it commits regarding transparency. Filled from public sources.

1 AGENTIC MATURITY SCALE

0

Classic automation

Fixed rules, no autonomy.

1

Assisted agents

The agent proposes, the human decides.

2

Governed agents

The agent acts within a frame, under supervision.

TARGET ~80%

3

Supervised autonomy

The agent decides alone; control after the fact.

2 YOUR KEY DECISIONS (filled)

A

Flagged assisted interaction

STAKE

Converse with the user while flagging as AI

IMPACT

Reversible, the user knows who they talk to

DEADLINE

Continuous

AUTONOMY

Autonomous (flags itself)

B

Marked content generation

STAKE

Produce text, image or audio marked as artificial

IMPACT

Traceable, origin disclosed

DEADLINE

Continuous

AUTONOMY

Supervised autonomy

C

Interact or release without disclosing the AI nature

STAKE

Let a human be assumed, release an unflagged deepfake

IMPACT

Misleads the user and commits the accountable party

DEADLINE

-

AUTONOMY

NON-DELEGABLE

Reading: transparency does not add a performance constraint; it sets, **before** deployment, what the agent must always disclose to the user.



OBJECTIVE

Define the agent's constitution: identity, standing principles and forbidden zones. Card filled with public case data.

1 AGENT IDENTITY & MISSION

AGENT NAME

Conversational agent or content generator (subject to art. 50)

ACCOUNTABLE (DDAO)

The provider or deployer of the AI system

PRIMARY MISSION

Interact with the public or produce content aimed at persons

SCOPE

Interaction and generation; under disclosure and marking obligations

2 CORE PRINCIPLES

- 1 Flag itself as AI at the start of the interaction, unless obvious
- 2 Mark all generated content in a machine-readable format
- 3 Disclose deepfakes as generated or manipulated
- 4 Disclose the origin of a public-information text

EXAMPLE PRINCIPLES

- Respect the legal and sector framework
- Protect data and confidentiality
- Stay transparent, every decision traceable
- When in doubt, escalate rather than decide

3 FORBIDDEN ZONES (ABSOLUTE)

- 1 Let the user ignore that they are talking to an AI
- 2 Release generated content without marking it as artificial
- 3 Publish a deepfake without disclosure

EXAMPLE PROHIBITIONS

- X Commit spending above a threshold without validation
- X Share personal data with a third party
- X Change its own rules or scope on its own

Constitution filled from the adopted text (EU AI Act, Article 50, obligations from 2 August 2026). It turns the legal transparency obligations into standing, signed rules.



OBJECTIVE

Continuously verify that transparency holds: flagging, marking, disclosure. Card filled with the publicly documented obligations of the case.

1 SUPERVISION INDICATORS

Indicator	Value / target (public sources)
Flagging of the AI nature	at the start of each interaction, unless obvious
Marking of generated content	machine-readable format, detectable
Disclosure of deepfakes	generated or manipulated, flagged
Public-information text	artificial origin disclosed
Oversight authority	national supervision, European framework

2 DEVICE

- 1 Verify, by sampling or automatically, that each interaction begins with the flag
- 2 Check that generated content carries its marking
- 3 Alert on failure: a coded rule can break silently
- 4 Renew the verification, not a box ticked once

An unsupervised transparency rule can break silently once in production. Linking supervision to the register (ACF-08) keeps the trace of disclosures.



OBJECTIVE

Keep the trace of the disclosures and markings made, to be able to account for them. Card filled with the fields relevant to the case.

1 FIELDS TO RECORD

Field to record	Example, AI Act art. 50 case
Timestamp	Date of the interaction or generation
Type	Conversational interaction / generated content
Disclosure made	AI flag / content marking
Marking format	Machine-readable, detectable
Status	Deepfake or information text: reinforced disclosure

2 USE OF THE REGISTER

- 1 Prove that the flag and marking did take place
- 2 Spot an undisclosed interaction or content
- 3 Establish who accounts for transparency: the provider or deployer
- 4 Provide verifiable proof of the artificial origin

Machine-readable marking is the verifiable trace of a content's origin. The register keeps its history and makes disclosure enforceable.

Card ACF-08 · Decisions register · completed with public case data.

Disclaimer

This study is carried out **exclusively from public information** (an adopted law and accessible legal analyses). It constitutes **neither an audit, nor a certification, nor legal advice**. It applies the ACF framework for pedagogical and illustrative purposes. No internal, confidential, or non-public data was used. The institutions cited belong to their respective owners.

Sources

- EU Artificial Intelligence Act, *Article 50: Transparency Obligations* - <https://artificialintelligenceact.eu/article/50/>
- European Commission, *Transparency obligations under Article 50 of the AI Act*
- Stibbe, legal analyses on the timeline and scope of article 50

ACF Case Study ACF-CS-023 · v1 · working document, to be reviewed before any public distribution.