

SERTEAM.

AI GOVERNANCE · FIXED-SCOPE ASSESSMENT

Your lawyer shouldn't be billing you to find your own AI systems.

Article 50 of the EU AI Act applies on 2 August 2026. The Digital Omnibus did not delay it. This document sets out exactly how the assessment runs, what it produces, and what it deliberately does not do.

Stalled

HOW THIS STARTS

Almost nobody buys this because of the deadline. They buy it because a deal stopped moving.

- 1 Which AI systems are you using?
- 2 Is AI-generated content disclosed to the people who see it?
- 3 Who is accountable for it?
- 4 Can you evidence any of it?

The answers usually exist. Nobody has ever assembled them. So the contract sits still while five people email each other about who owns the chatbot – and the 2 August deadline is about to make these questionnaires far more common, not less.

Moved

WHAT MOVED, WHAT DIDN'T

The Omnibus moved two deadlines. Article 50 was not one of them.

Delayed

Annex III standalone high-risk systems: 2 August 2026 → 2 December 2027. Annex I embedded high-risk: 2 August 2027 → 2 August 2028.

Not moved

Article 50 transparency duties apply from 2 August 2026, exactly as legislated. One narrow exception: machine-readable marking under Art 50(2) has grace to 2 December 2026, and only for systems already placed on the market before August. It covers no deployer obligation.

Role

WHAT THE ASSESSMENT CLASSIFIES

Four obligations. Two roles. The classification decides everything.

50(1) · 50(2)

Provider duties. Systems interacting directly with people must make clear a person is dealing with AI. Synthetic audio, image, video and text must be marked machine-readably.

50(3) · 50(4)

Deployer duties. Inform people exposed to emotion recognition or biometric categorisation. Disclose deep fakes, and AI-generated text published to inform the public on matters of public interest.

Art 3(3)

The pivot. You are a provider where a system is placed on the market under your own name or trademark. Plain use of an off-the-shelf tool is not that. A commissioned build or a white-labelled product is — which moves 50(1) and 50(2) onto you. This is the most common finding.

Proof

THE 50(4) DEFENCE

You probably have the defence. You probably can't prove it.

Record.

Article 50(4) exempts AI-generated text published to inform the public where there has been human review or editorial control, and a natural or legal person holds editorial responsibility.

The test is disjunctive — review or control, not both. That is more forgiving than most firms assume. But the floor is real: the Commission has been explicit that superficial, solely formal or procedural checks, such as spell-checking or grammar, do not amount to review or editorial control.

The typical mid-market position is that the editorial process genuinely exists and there is no record of it. No named person holding responsibility. No log of what was reviewed, by whom, or when. Day 4 tests exactly that, and the output is a list of what is missing.

Five

HOW THE ENGAGEMENT RUNS

5

working days of billed effort, delivered across a two-week window to fit around your people's diaries.

Days 1–2

Scope, sweep and role classification. Discovery across procurement records, SaaS spend and the marketing, product and service stack — including AI embedded in tools you already bought. Four to six 45-minute interviews. Every system tested against Art 3(3) and 3(4).

Days 3–4

Duty mapping and the defence test. Each system mapped against 50(1)–(4) and the timing requirements of 50(5). Disclosures reviewed as a user meets them. Residual legal questions drafted with the facts attached.

Day 5

Plan and readout. Remediation sequenced and owner-assigned, split into before 2 August, before 2 December, and out of scope. Board one-pager. 45-minute readout to your exec team.

SERTEAM.

€4,500 · FIVE WORKING DAYS

What you're left holding.

- 1 AI Content & Interaction Register — every in-scope system, with owner, procurement route and role classification.
- 2 Article 50 Duty Map — system by system, sub-article by sub-article, RAG-rated, with the reasoning shown.
- 3 Editorial Control Defence Assessment — whether 50(4) holds in practice, and exactly what is missing if it does not.
- 4 Counsel Question Set — the remaining legal questions, framed with the facts attached.
- 5 Disclosure & Marking Specification — wording for user-facing disclosures and the technical approach to marking.
- 6 30/60/90 Remediation Plan — sequenced, owner-assigned, deadline-split, plus a one-page board summary.

What you provide: a named executive sponsor, access to procurement records and your vendor list, four to six 45-minute interviews, and contracts for AI-enabled tools.

BOOK A SCOPING CALL →

This assessment is not legal advice and does not constitute a legal opinion. SER Team is not a law firm and does not hold itself out as one. Findings are provided to inform your own decisions and to support review by qualified legal counsel. Total liability is limited to the fees paid for the engagement. · Sources: Reg (EU) 2026/1744 · European Commission, Article 50 transparency FAQ · Code of Practice on Transparency of AI-Generated Content (10 June 2026)