

Specimen document. This is a real Evidence Pack rendered from a demonstration scan, so you can see the format before buying one. The hashes and seal below are genuine for that scan and verify normally; the site it describes is a fixture, not a customer.

ACME COMPLIANCE GMBH

AI-disclosure scan record

EU AI Act, Article 50: technical findings with sealed evidence

TARGET	https://demo-fixture.example
SCAN ID	aaaaaaaa-aaaa-4aaa-8aaa-aaaaaaaaaaaa
CAPTURED	2026-07-27 01:01:55 UTC → 2026-07-27 01:03:55 UTC
RULE PACK(S)	eu-ai-act-art50 2026.07.7
CAPTURE METHOD	Cloudflare Browser Rendering (homepage) + capped static fetch (crawled pages, when present)
EVIDENCE MANIFEST SHA-256	f1e2d3c4b5a69788f1e2d3c4b5a69788f1e2d3c4b5a69788f1e2d3c4b5a69788

TAMPER-EVIDENT: every artifact in this pack is SHA-256 hashed into the evidence manifest; the manifest hash above is HMAC-signed server-side at seal time. Independent verification: recompute any artifact hash and compare against the manifest (appendix).

Findings

States are technical indicators (detected / not detected / could not verify) graded under the versioned rule pack above, not legal conclusions.

ATTENTION Art. 50(1)

Chatbot AI disclosure at first interaction

No AI/automation disclosure was detected in the chat widget's first interaction (vendor: intercom). If this assistant is AI-driven, Article 50(1) requires disclosing that to users.

Root cause: the chat panel's first interaction was read successfully and contains no AI/automation disclosure term. Automation could not be confirmed by the reply-latency probe, so this stays a hedged WARN rather than a FAIL — but if this assistant is AI-driven, Art. 50(1) applies and the greeting or first automated reply needs an explicit AI disclosure.



Captured 2026-07-27 01:02:55 UTC · SHA-256 4444555566667777888899990000bbbb4444555566667777888899990000bbbb

How to fix — Intercom (Fin)

1. Open your site in a private window as an EU visitor and start a chat: is it explicit you're talking to AI before or in the very first exchange?
2. Check Fin's identity settings, keep the AI descriptor visible and don't rename Fin to a human-sounding persona without an AI label
3. Put an explicit line in the opening message (e.g. "I'm Fin, an AI agent, I can bring in a human anytime") rather than relying on an icon
4. Verify the disclosure survives on mobile and in the compact Messenger view, not just desktop
5. Check the human-handover message makes the participant change clear in both directions
6. Screenshot the first-interaction state and save your Messenger/Fin settings export with a date

Known failure mode: The classic Intercom failure mode is cosmetic humanization: Fin renamed to a teammate-style persona with a human avatar, answering instantly with natural language. Each choice is individually reasonable; together they make the AI nature non-obvious, which is exactly what strips you of the 'obvious from context' exception. The second one is drift, Fin ships new capabilities and settings frequently, and a Messenger update or workspace migration can quietly change what visitors see first.

ATTENTION Art. 50(2)

Machine-readable provenance marking on media

AI-generation indicators were found on 1 of 2 sampled assets, and 1 sampled asset(s) carry no machine-readable provenance marking (C2PA/IPTC). Under Article 50(2), AI-generated content must carry such marking where technically feasible.

Root cause: this site demonstrably publishes AI-generated media (1 of 2 sampled assets carry AI-generation indicators), and 1 sampled asset(s) carry no machine-readable provenance marking at all. We cannot tell from outside which of those unmarked assets were machine-generated, so this is flagged for review rather than reported as a gap. The usual cause is a CMS or CDN image-optimization step that strips C2PA/XMP at publish time: the generator attached credentials, the pipeline removed them, and removing them also removes the indicator that would let anyone tell which assets were affected.

Sampled assets: 2 checked · 1 with marking

93.184.216.34/assets/hero-illustration.png	xmp · AI-indicated · adobe firefly
93.184.216.34/assets/blog-cover.jpg	no marking

How to fix

1. For AI-generated media, keep the C2PA Content Credentials your generator embeds (DALL-E, Adobe Firefly, and others attach them by default) — many CMS/CDN image pipelines strip metadata on optimization; check yours preserves it
2. If your pipeline strips metadata, re-attach credentials at publish time (c2patool or the CAI SDKs) or add IPTC DigitalSourceType=trainedAlgorithmicMedia
3. Media that is NOT AI-generated needs no marking — Article 50(2) applies to AI-generated content only
4. Keep an export of your pipeline configuration with a date, so you can show when marking was in place

ATTENTION Art. 50(2) (supporting)

Provenance marking is structural, not textual-only

Provenance markers on the sampled assets are textual mentions only; no structural C2PA manifest was detected on any of 1 marked asset(s).

Root cause: provenance markers were found on 1 sampled asset(s), but none carried a structural C2PA (Content Credentials) manifest store — the markers are textual mentions (IPTC fields or generator strings) only. A structural manifest is the robust machine-readable form: it binds the claim to the asset and survives scrutiny where a text field does not. The usual cause is a pipeline that re-encodes assets and drops the embedded JUMBF box while leaving lightweight metadata behind.

ATTENTION Art. 50(4)

AI-content label on article-like pages

One signal, uncorroborated: 1 independent kind(s) of publisher-emitted signal point to AI use in content production, and no AI-content label was detected on 3 of 4 article-like page(s) (6 page(s) scanned). This is evidence about the site, not a determination about any page's text. Manual verification recommended.

Root cause: one publisher-emitted signal indicates AI tooling is in use here, uncorroborated by any other kind of evidence, and 3 of 4 article-like page(s) carry no AI disclosure. A single signal does not establish which pages a tool produced. Confirm which content the tool was used for.

One signal, uncorroborated 1 independent kind of signal · 1 observation

Signals the site itself emitted about using AI to produce content. These are the publisher's own declarations, recorded as found; DisclosureProof does not assess whether any text was machine-written and reports no likelihood of that.

MEDIA PROVENANCE	1 of 2 sampled asset(s) carry AI-generation provenance site-wide
------------------	--

Affected pages: 3 unlabeled article page(s)

<https://demo-fixture.example/blog/ai-in-logistics>

<https://demo-fixture.example/blog/q3-market-outlook>

<https://demo-fixture.example/news/expansion-2026>

How to fix

1. If any published text is AI-generated on a matter of public interest, add a visible label (the Code of Practice's uniform wording/icons are the safe default)
2. Human-reviewed content with editorial responsibility can rely on the Art. 50(4) carve-out — document the review process
3. Add the label near the byline or article footer where a reader actually sees it, not only in metadata
4. State your AI-use policy on a dedicated page and link it site-wide

PASS Art. 50 (supporting)

AI policy / disclosure page

A dedicated AI policy/disclosure page was found (<https://demo-fixture.example/ai-policy>). This is supporting context that helps demonstrate transparency.

N/A Art. 50(1)

Chatbot disclosure persists on mobile

Mobile disclosure was not separately assessed (no desktop disclosure to compare, or the mobile panel could not be read).

N/A Art. 50(1)

Chatbot disclosure covers the site's languages

Language coverage was not assessed for this scan — it requires a detected chat widget, a detected disclosure, and readable site-language markup (html lang / hreflang).

N/A Art. 50(4) (supporting)

Official EU labelling icons in use

Not assessed: applies only when article-like pages already carry an AI-content label (4 article-like page(s) found).

N/A Art. 50(3)

Emotion recognition / biometric categorisation disclosure

Article 50(3) (emotion recognition and biometric categorisation systems) cannot be verified by an external crawler. Complete the in-product attestation module to cover this duty; attested answers are sealed as evidence like any other artifact.

Article 50(4) attestation (self-attested)

Publishes AI-generated text	Yes
Human-reviewed under editorial responsibility	No
Attested by	Maxwell Milan (Editor), editor@example.com
Attested at	2026-07-25T16:25:14.308Z
Scan observed at attestation time	UNVERIFIED

The answers above were provided by the operator (self-attested). DisclosureProof records and seals them as evidence of the operator's declaration, and combines that declaration with what the scan itself observed about labelling. It does not and cannot independently determine whether any text was machine-generated.

Attestation record SHA-256: fa82c9e14a732dc6d61d168ad63b3fbf8a21d60fe52e01e2b572c4f4e64e9d05 (HMAC-signed)

Appendix: sealed artifact hashes

Every artifact captured for this scan, as recorded in the evidence manifest at seal time (—, crawler v?).

ARTIFACT	SHA-256	BYTES