

Risk Classification Memo

A signed, hash-chained evidence record — see the final section for how to verify it independently, without relying on Clause50.

41.7%

Coverage score

2 of 6

Requirements satisfied

4

Active items

5

Evidence items



System	Demo — Campaign image generator
System ID	cmu6ekeln002kga1khhim53ed
Organisation	Sablewick Studio (specimen)
Purpose	Internal tool that turns a campaign brief into finished campaign images. A brief assistant converses with account managers and designers to draft image directions; a generation pipeline renders photorealistic images, including scenes with people and places, for client review. Built and operated by Sablewick Studio on a licensed image model.
Actor role	Both
Rule pack	2026.19
Template	Risk Classification Memo
Generated	18 September 2026, 03:32 UTC (2026-09-18T03:32:24.329Z)

Classification

What the recorded answers engage: Article 50 transparency — applies now

3 obligation(s) in this rule pack apply to this system on its recorded answers, all of them in force today.

Not confirmed. Nobody at the organisation has reviewed or confirmed this reading. It is printed as unconfirmed rather than omitted: what was recorded, and what was not, are both facts about this system's record.

This states the scope of the evidence Clause50 is assembling under the rule pack cited above, read off the answers recorded for this system. It is not a legal classification of this system, and it is not legal advice. Any confirmation shown above is the operator's own statement, recorded and reproduced by Clause50, not assessed by it. Put this document to your legal adviser.

1. System facts

Fact	Recorded value
System name	Demo — Campaign image generator
System ID	cmu6ekeln002kga1khhim53ed
Organisation	Sablewick Studio (specimen)
Intended purpose	Internal tool that turns a campaign brief into finished campaign images. A brief assistant converses with account managers and designers to draft image directions; a generation pipeline renders photorealistic images, including scenes with people and places, for client review. Built and operated by Sablewick Studio on a licensed image model.
Deployment context	Sablewick Studio (fictional) — a marketing agency with offices in Dublin and Lisbon. Creative Technology built and runs the tool; the studio team uses it daily. Still a pilot at version 0.9; a v1.0 release is planned once output marking is complete.
Actor role	Both
Current system version	0.9.3
Model provider	— (not recorded)
Model name	— (not recorded)
Model version	— (not recorded)
Configuration hash	— (not recorded)
Source commit	— (not recorded)

Every value above is reproduced from this system's own record and its current version as held at the moment this document was generated. Nothing in this section is inferred, derived, or assessed.

2. Annex III screening

Screened: none recorded. All 8 questions below were answered for this system, and none was answered yes. This section records the screen itself — it is not an absence of evidence.

Question	Recorded answer
Does the system identify people, sort them into categories, or infer their emotions using biometric data?	No
Is the system a safety component in the operation of critical infrastructure?	No
Is the system used to decide access to education, to evaluate learning, or to monitor students during assessments?	No
Is the system used in recruitment, or to make or support decisions about people at work?	No
Is the system used to decide whether people get access to essential public or private services?	No
Is the system used by, or on behalf of, law enforcement?	No
Is the system used in migration, asylum, or border-control management?	No
Is the system used to assist judicial decision-making, or to influence elections or voting behaviour?	No

Every answer above is reproduced from this system's own record as held at the moment this document was generated. Nothing in this section is inferred, derived, or assessed.

3. Art. 5 prohibited-practice screening

Screened: none recorded. All 8 questions below were answered for this system, and none was answered yes. This section records the screen itself — it is not an absence of evidence.

Question	Recorded answer
Does the system use subliminal, purposefully manipulative, or deceptive techniques that change how people behave?	No
Does the system target people's age, disability, or social or economic situation in order to change their behaviour?	No

Question	Recorded answer
Does the system score or rank people based on their social behaviour or personal characteristics?	No
Does the system predict how likely a person is to commit a criminal offence?	No
Does the system build or expand a facial-recognition database by untargeted scraping of images?	No
Does the system infer people's emotions in a workplace or an education setting?	No
Does the system use biometric data to deduce race, political opinions, trade-union membership, religious or philosophical beliefs, sex life, or sexual orientation?	No
Does the system identify people in real time, at a distance, in publicly accessible spaces?	No

Every answer above is reproduced from this system's own record as held at the moment this document was generated. Nothing in this section is inferred, derived, or assessed.

4. Art. 50 transparency triggers

Deployers of an AI system that generates or manipulates image, audio or video content constituting a deep fake shall disclose that the content has been artificially generated or manipulated. Deployers of an AI system that generates or manipulates text which is published with the purpose of informing the public on matters of public interest shall disclose that the text has been artificially generated or manipulated, save where the AI-generated content has undergone a process of human review or editorial control and a natural or legal person holds editorial responsibility for publication of the content.

Providers shall ensure that AI systems intended to interact directly with natural persons are designed and developed in such a way that the natural persons concerned are informed that they are interacting with an AI system, unless this is obvious from the point of view of a natural person who is reasonably well-informed, observant and circumspect.

Providers of AI systems, including general-purpose AI systems, generating synthetic audio, image, video or text content, shall ensure the outputs of the AI system are marked in a machine-readable format and detectable as artificially generated or manipulated. (Code of Practice, June 2026, is the practical benchmark. Systems on market before 2026-08-02 have until 2026-12-02 per Omnibus.)

Art.50(4) — Deepfake disclosure

STATUS: SATISFIED

Labelling/disclosure mechanism for deep-fake image, audio or video outputs (visible label, provenance metadata, or C2PA manifest per modality), uploaded and attested item by item against the checklist.

Proven by: verified marking manifest (T3) — generated by Sablewick Campaign Image Generator/0.9.3 c2pa-rs/0.36.0 and signing certificate issued by Sablewick Studio Content Credentials (demo).

Clause50 parsed this manifest and recorded its claim generator, its assertions and its signing certificate's issuer. **Clause50 did not independently verify the signature** and makes no statement about its cryptographic validity.

This evidence establishes that the artefacts submitted were marked. **It does not establish that every output of the system is marked** — no sample was declared with it.

Art.50(1) — Disclosure implementation

STATUS: PARTIAL

Screenshot or config showing AI-disclosure at or before first interaction (UI text, system prompt, or API metadata), uploaded and attested item by item against the checklist.

Art.50(1) — Disclosure in traces

STATUS: MISSING

Sampled conversation traces demonstrating disclosure precedes first substantive turn.

Art.50(3) — Emotion biometric notice

STATUS: NOT APPLICABLE

Notice shown to exposed persons disclosing operation of the emotion-recognition / biometric-categorisation system (UI text, signage, or in-product notice), uploaded and attested item by item against the checklist.

Art.50(2) — Marking mechanism

STATUS: MISSING

Watermark/metadata/provenance mechanism per applicable modality (C2PA, metadata tags, standardised EU label), uploaded and attested item by item against the checklist.

Proven by: operator attestation. No machine-readable marking artefact was submitted for this requirement, so the statement above rests on the operator's checklist rather than on evidence Clause50 could read.

Art.50(4) — Public text disclosure

STATUS: MISSING

Disclosure mechanism for AI-generated text published to inform the public on matters of public interest (byline/label or editorial-responsibility record), uploaded and attested item by item against the checklist.

Art.50(1) — System in operation

STATUS: SATISFIED

System in active production use during the period — machine-observed usage activity (request/token counts) from a connected provider when available, or an operator attestation otherwise.

5. Conclusion and legal basis

What Clause50 is preparing evidence for.

In force now:

- **Art.50(1)** — Disclosure of AI interaction

- **Art.50(2)** — Machine-readable marking of synthetic content

SPECIMEN — Clause50 demo document. Not evidence of any real AI system.

Generated by Clause50 · for Sablewick Studio (specimen) · clause50.com ·
Verify: <https://clause50.com/verify>


- **Art.50(4)** — Disclosure of deep fakes and AI-generated public-interest text

On what basis.

16 screening question(s) were put to this system in this document. 0 were answered yes; 0 were not answered. The full questions and answers are reproduced in the screening section(s) above.

This section states the scope of the evidence Clause50 is assembling under this rule pack. It is not a legal classification of this system, and it is not legal advice. Clause50 records what was answered and which obligations follow from the rule pack cited above; it does not assess whether an obligation applies to you in law, and the underlying obligations carry express derogations this service does not evaluate. Put this document to your legal adviser.

6. Cryptographic signature seal



This document is signed, auditor-readable evidence of what your system can demonstrate about its compliance at the time of generation.

What is signed:

- Per-requirement status
- Per-item evidence hash (immutable audit trail of what was submitted)
- Per-system chain-head hash (append-only; any tampering is detected)
- An Ed25519 detached signature over the manifest

Key: key-specimen-test

Chain head: 0310ba26918bac8302a6a3b7ab4a540ac56dd03e191ff7880e35bbb9dfe6c1ec

Evidence items: 5

Rule pack: 2026.19

How to verify:

1. **Online:** Drop the ZIP at <https://clause50.com/verify>
2. **CLI:** `npm run verify -- manifest-bundle.json --public-key <key>` (key from the registry)
3. **Manual:** Extract the ZIP, curl the registry, and use tweetnacl to verify the signature

7. Disclaimers and scope limitations

7.1 What this proves (in scope):

- This document was produced by Clause50 and has not been altered since it was signed.

7.2 What this does NOT prove (out of scope):

- That the evidence is accurate or complete.
- That your system actually complies with the regulation (compliance is an operational and legal matter, not a document matter).
- That you have submitted all the evidence you have. A party holding both this database and the signing key can produce a valid manifest over a curated subset of evidence. This is a limit every compliance-audit tool has, and most do not admit it.
- That an evidence item's CONTENT matches what it claims to be. Clause50 confirms an item's presence, timestamp, and cryptographic hash, and — for a manual attestation — the checklist the operator claimed; it does not independently parse or verify that an uploaded watermark, disclosure notice, or marking mechanism actually functions as described.
- Section 6a lists the evidence this document's scoring actually relied on for each requirement; section 6b lists rolled-up event-kind activity; section 6c lists evidence held in your chain and shipped with this pack that no requirement in this pack relied on.
- Evidence images reproduced in this document are resized and re-compressed for print; they are re-verified against the recorded content hash before rendering, but the embedded picture is not byte-identical to the uploaded file. The original, full-resolution file is included unmodified in this pack's export ZIP, under `evidence/`.