

The Producer's **AI** Compliance Checklist

The compliance and documentation a production should have in place whenever AI is part of the work, across film, television and commercial production.

Four questions every producer should be able to answer

If AI was used on a production, a producer should be able to answer four questions:

- Which AI tools were used?
- What content did they affect?
- Who approved their use?
- Where is the record?

Most compliance, rights and governance questions reduce to those four points. If you can answer them clearly, you are usually in a far stronger position than if you are reconstructing events months later.

This checklist covers four areas: legal obligations, copyright and rights management, tool governance, and production oversight. A dated record of AI use supports all four.

The details of AI regulation will continue to evolve. The value of keeping a clear production record does not.

Copyright and rights

Copyright in AI-assisted work depends on demonstrable human creative input. Your ability to reserve rights over your own material depends on knowing what content exists and where it is distributed. Both depend on a record.

PARTS 3 · 4

Tool governance

Which AI tools were used, what material was entered into them, and who approved their use. The difference between governed use and unapproved use is often the difference between a manageable issue and an expensive one.

PARTS 5 · 7

Production record

A dated account of how AI was actually used across the production. Useful for legal, contractual and operational questions long after delivery.

PART 8

Regulatory obligations

The EU AI Act introduces disclosure and AI-literacy obligations. Meeting them starts with knowing what tools were used and where their outputs appeared.

EU · PARTS 1 · 2 · 6

How to use this checklist. It works through each of these in turn. Tick what is in place, and use the gaps as your to-do list. Written for producers working in or distributing into the EU, with notes where the UK position differs. General information current as of June 2026, not legal advice.

1 Transparency and labelling (Article 50)

2 Aug 2026

Transparency rules apply (Art. 50)

Deployers (50(4))

Disclose deepfakes; no editing carve-out

Penalties

Up to €15M or 3% turnover

The transparency rules apply from **2 August 2026**. The duty that matters most to producers is the **deployer's (Article 50(4))**: if you use a generative AI system to create a deepfake, you must disclose that the content is artificially generated, clearly and at first exposure. The Commission's draft guidance (8 May 2026) confirms this applies **even without intent to deceive**. Ordinary post-production such as colour grading or object removal does not by itself create a deepfake. Note that the standard-editing carve-out sits in the provider marking duty (50(2)), not your disclosure duty (50(4)), which turns only on whether the output meets the deepfake definition in Article 3(60). Generative fill or AI frame-extension, which create content rather than tidy it, can trigger disclosure even where provider marking would not. Fictional and artistic works keep a lighter form, adapted so disclosure does not spoil the work, unless the content is primarily informative or commercial. The Code of Practice on marking and labelling (10 June 2026) is voluntary but is becoming the practical benchmark. Breach can reach €15M or 3% of worldwide turnover. A proposed simplification (the AI Omnibus, not yet adopted) would briefly delay the 50(2) marking duty for systems already on the market, but not the 50(4) duty. **UK note:** no direct UK equivalent, but content distributed into the EU is in scope.

- ☐ Map where AI-generated or AI-manipulated content reaches an audience: final cut, trailer, marketing, social.
- ☐ Identify any **synthetic version of a real, identifiable person** and plan clear, first-exposure disclosure for it.
- ☐ For AI-generated text published to inform the public, check whether the editorial-oversight carve-out genuinely applies. A light human "check" is not enough.
- ☐ Put a labelling and marking clause in vendor and post-house agreements so provenance is supported upstream.
- ☐ Decide whether to sign or to align your practice with the Code of Practice.

2 AI literacy (Article 4)

The AI literacy duty has applied since **February 2025**; national authorities gain enforcement powers from **2 August 2026**. It is proportionate and role-based: staff and contractors who use AI need an understanding appropriate to their role, not a certificate. A proposed amendment (the Digital Omnibus, politically agreed May 2026) would soften the wording but is **not yet law**, so the current duty stands. The penalty tier sits at €7.5M or 1% of turnover.

- ☐ Confirm the people using AI on your productions understand what the tools do and where they fall short.
- ☐ Keep a simple record of any briefing or training, tied to the tools actually in use.
- ☐ Treat this as foundational: the same inventory feeds every other obligation here.

3 Copyright and authorship

Copyright protection depends on human creative contribution. Where a person directs, selects, edits and shapes AI-assisted output, the claim to authorship is generally stronger than where material is generated with minimal human involvement.

In practice, productions rarely contain neatly separated categories of "AI-assisted" and "AI-generated" material. Assets often pass through multiple tools and iterations. What matters is not attaching a permanent label to every output, but preserving evidence of the human creative process behind it.

No law requires a producer to maintain that record. However, a producer who can show how decisions were made is in a stronger position than one relying on recollection after the fact.

- ☐ For AI-assisted material in the final work, record the human direction, selection, iteration and editing behind it.
- ☐ Retain the decision trail, prompts included, wherever it is practical to do so.
- ☐ Treat the **assisted**-versus-**generated** split as descriptive, not a clean per-asset stamp; most clips mix both across many iterations.

4 Reserving your rights (TDM opt-out)

Under Article 4 of the DSM Directive, commercial text and data mining of your content is generally permitted unless rights have been expressly reserved through an appropriate machine-readable mechanism.

For audiovisual works, this responsibility typically sits with the producer. Unlike music, there is no sector-wide rights-management structure making that decision on your behalf. If rights have not been reserved, your content may be available for lawful text and data mining.

- ☐ Decide your position on text and data mining of your catalogue and back catalogue.
- ☐ Reserve rights by an appropriate machine-readable means wherever you control distribution.
- ☐ Treat opt-out as the producer's decision, not the platform's and not a collecting society's.

5 Data handling and confidentiality

AI tools differ significantly in how they store, retain and use submitted material. Scripts, budgets, contracts, development documents and unreleased footage may all create different risks if they leave your control.

The issue is not only confidentiality. It is also control over commercially sensitive production information. Before sensitive material is entered into any AI system, understand how that system handles retention, training and access.

- ☐ Define what may and may not be put into AI tools, by material type.
- ☐ Check each tool's terms for training and retention before anything sensitive goes in.
- ☐ Prefer enterprise or no-training settings for confidential material.

6 Talent, likeness and synthetic performers

Synthetic representations of real people are likely to attract the greatest scrutiny on most productions. They raise questions of consent, contractual rights, reputation and disclosure. Under Article 50, certain uses may also trigger transparency obligations. Because these issues often overlap, any use of digital replicas, synthetic voices or altered likenesses should receive specific review before release.

- ☐ Get specific written consent for any digital replica, synthetic voice or altered likeness.
- ☐ Check what talent contracts, guild terms and commissioner agreements require.
- ☐ Flag every synthetic-likeness use for a labelling assessment.
- ☐ Review AI-generated people and places for accidental likeness to real, identifiable persons before the final cut. Resemblance can trigger disclosure with no intent, and the production record lets you trace where any such output entered.

7 Crew and vendor governance

Many productions have formal policies for AI use. Fewer know whether those policies are actually being followed.

The most common governance gap is unapproved tool use by freelancers, vendors or individual crew members. Once AI-generated material enters the production pipeline, identifying it after the fact can be difficult. Clear contractual requirements reduce that risk.

- ☐ State which AI tools are approved and which are off-limits in deal memos and supplier contracts.
- ☐ Ask crew and freelancers to disclose the AI tools they used.
- ☐ Record the licence tier each tool ran on, not just its name. The same tool on a free versus enterprise plan carries different training and retention terms.
- ☐ Name one person responsible for the production's AI record.

THE SHORT VERSION

Four questions a producer should be able to answer

- | | |
|-----------------------------|---------------------------------|
| ① Which AI tools were used? | ② What content did they affect? |
| ③ Who approved their use? | ④ Where is the record? |

Keep the record

Every section in this checklist ultimately depends on the same thing: a dated record of how AI was used during production.

Questions about authorship, disclosure, contractual compliance and provenance often arise long after a project has been delivered. By that point, memory is unreliable, freelancers may have moved on, and project files may be dispersed across multiple vendors. A production record provides a single reference point.

The objective is not to prove compliance with every rule. It is to preserve an accurate account of what happened.

Keep one record per production.

Record:

- the AI tool used
- the person using it
- the purpose of the use
- whether the output entered the final work

Store the record in a format that remains accessible after delivery and independent of any individual team member.

Audiovisual AI Log creates that record

Music has long had its cue sheet, and the camera department its shot list. AI use has had no equivalent.

Crew complete a short, structured survey, and responses compile into a single dated record of AI use across the production.

Each entry is classified by how the output was used:

TIER 1

Not in final

AI used in support; output did not reach the cut.

TIER 2

AI-assisted

Human-directed output that entered the work.

TIER 3

AI-generated

Generated output that entered the work.

The record exports in multiple formats, including a formatted PDF for review and structured JSON for other systems. Because every entry is structured, the data also supports analysis across a production: the volume and distribution of AI use, and indicative cost or carbon estimates where the data allows.

The result is one dated record a producer can hand to a broadcaster, an insurer, a co-producer or legal counsel, rather than reconstructing the production after the fact.

AI documentation is becoming part of normal production governance

We are speaking with producers, production companies and commissioners across Europe about how these records should work in practice and what information is genuinely useful to keep.

If you run productions that use AI and want a structured, dependable way to document it, we would welcome the conversation. We are working directly with a small number of productions to shape the tool around real use.

[Start a conversation →](#)[Read more →](#)jaakko@availog.com

Sources. Regulation (EU) 2024/1689 (EU AI Act), Articles 4 and 50; European Commission draft guidelines on Article 50 (8 May 2026); Code of Practice on marking and labelling of AI-generated content (10 June 2026); Digital Omnibus political agreement (7 May 2026); Directive (EU) 2019/790 (DSM), Article 4. General information, current as of June 2026. Not legal advice and does not create a lawyer-client relationship; obligations vary by role, content and jurisdiction.