




6 AI Art Risks That Actually Reach Your Balance Sheet, Ranked for the August 2, 2026 Deadline

“Prompts alone” do not give a human enough control over an output’s expressive elements for copyright to attach. That is the US Copyright Office’s January 2025 position, and as of August 2026 no reported US federal decision has displaced it.

WHERE THIS LANDS

The risk that costs you money first is the EU labeling duty: Article 50 of Regulation (EU) 2024/1689 applies generally from August 2, 2026, and it binds both providers who generate synthetic output and deployers who publish deepfakes. The ownership gap from the Copyright Office’s January 29, 2025 Part 2 report is slower, cheaper to fix, and almost never covered by a vendor indemnity.



Six exposures, ranked by how fast they hit a Q4 invoice

I am ranking six exposures a creative or product team carries when it ships AI-generated images, audio, video or text into a Q4 2026 campaign. The test is practical: how fast does this cost real money, and can you buy your way out of it with a vendor contract?

That second half matters more than teams expect. Vendor “commercially safe” marketing is overwhelmingly a copyright indemnity. Copyright is one of six items below, and on my reading it is not the one most likely to land on a Q4 invoice.

Missing machine-readable marking on generated output

Article 50 requires providers of AI systems generating synthetic audio, image, video or text to mark outputs in a machine-readable format and make them detectable as artificially generated or manipulated. The duty applies generally from August 2, 2026, and no indemnity transfers it.

Undisclosed deepfakes in a published campaign

Article 50 puts a separate duty on deployers of systems producing deepfake image, audio or video to disclose that the content is artificially generated or manipulated, so the agency that publishes is exposed even when the model vendor did its part.

Owning nothing you thought you owned

The Part 2 report concludes prompts alone generally do not confer copyright, so an asset library built on prompting may have no exclusive rights behind it and no basis to stop a competitor reusing the look.



Trademark and likeness in the generated frame

Getty's claim against Stability AI, IL-2023-000007, bundles trademark alongside copyright, and the England and Wales High Court's approved judgment of January 14, 2025 did not finally determine those substantive claims. The exposure sits outside what a copyright-only indemnity reaches.

Training-data liability flowing back to buyers

Thomson Reuters won a \$25 million jury verdict against Ross Intelligence in the District of Delaware on February 11, 2025, and Bartz v. Anthropic produced a reported \$1.5 billion settlement over books allegedly used in training. The upstream numbers are large, even though these are provider-side cases.

Public-interest text published without disclosure

Deployers publishing AI-generated text on matters of public interest must disclose it unless the content went through human review and a person or organisation holds editorial responsibility. That is a narrow exemption and an easy one to document.

The order is mine. Each line rests on the sourced facts cited in it; the sequencing reflects my judgment about speed to balance sheet, not any published enforcement ranking.

Items one and two rank above the ownership gap for a simple reason. They are live obligations with a date attached, and they attach to the act of publishing. The ownership gap is a loss of upside. You cannot sue the person who copies your unprotected image, which hurts, but it does not produce a demand letter in November.

Vendor "commercially safe" promises are written around copyright claims. Trademark,



personality and likeness, and the Article 50 marking and disclosure duties are separate legal tracks. **Read the indemnity clause for what it names, not for what the landing page implies.**

No, the exception for assistive editing covers more than you think

The European Commission's [quick facts on transparency rules](#) recognise an exception where the AI system performs an assistive editing function or does not substantially alter the input data or its semantics. Generative fill that repairs a sensor dust spot is a different thing from a generated hero image, and the text of [Regulation \(EU\) 2024/1689](#) treats them differently.

The practical problem is that the exception is defined by substance, not by tool. Your retouchers decide, per asset, whether the semantics changed. That decision needs to be recorded somewhere an auditor can read it, and most asset pipelines have no field for it.

Marking is also only half the chain. I have written before about [how Content Credentials actually survive distribution](#), and the short version is that embedded provenance does not reliably reach the viewer. The regulation asks providers to mark. It does not make the platforms in between preserve the mark.

Part 2 protects the human layer, not the prompt

The Part 2 report is frequently read as "AI output cannot be registered." That is wrong. The [report](#) names protectable human contributions: perceptible human-authored expression inside the work, creative selection, coordination and arrangement of AI material, and creative modification of AI output. Copyright Office materials from January 2025 identify more than 1,000 registrations containing AI-generated material, with the AI portions disclaimed where required.

So the asset is not worthless. The protection sits in the human layer, and only if someone can show the human layer existed. Composition decisions, overpainting, the arrangement of twelve generated elements into one key visual: that is where the claim lives, and it needs



to be evidenced at the time rather than reconstructed later.

What I left off: the doomsday scenarios and the Anthropic headline

I left out “the models will get regulated out of existence” and “the courts will void everything.” Neither is supported. [Chambers’ August 12, 2026 USA survey](#) states that no reported US federal decision had displaced the Copyright Office’s registration approach, and the Getty judgment of January 14, 2025 left the substantive copyright and trademark questions open rather than resolving them either way. The pending uncertainty is real, but it sits upstream, between rightsholders and model providers. It is not the thing that reaches your Q4 close.

I also left out the \$1.5 billion Anthropic figure as a planning input for buyers. It is a reported settlement subject to court approval, not a merits damages award, and settlements price litigation risk and appetite as much as liability.

MY TAKE

The two August 2, 2026 duties will be enforced through procurement and platform policy long before any regulator writes a decision. A brand’s EU agency will start asking for marking attestations because its own counsel told it to, and vendors without an answer will quietly lose work. That is a commercial mechanism rather than a legal one, and it moves faster. I expect the ownership question to stay unresolved in US courts well into 2027, which means the Part 2 framework (human expression, arrangement, modification) is the operating assumption for at least the next planning cycle. I could be wrong about the timing; one appellate decision would change the picture.


The cheapest control across all six items is the same control: a per-asset record of what the model produced, what a human changed, and who signed off on whether the semantics changed. That record answers the Article 50 marking question, supports the deployer disclosure, and documents the human contribution the Copyright Office asks for. **One field, three problems.**

What it does not solve is trademark and likeness. Nothing in the record tells you whether



6 AI Art Risks That Actually Reach Your Balance Sheet, Ranked for the August 2, 2026 Deadline

the generated logo on the generated shopfront belongs to someone. That still requires a human looking at the frame before it ships, and the more than 10,000 public comments the Copyright Office reviewed did not produce a tool that does it for you.



WRITTEN BY

Artur Markus

Software engineer in Basel building AI infrastructure and agentic systems, with a background in pharmaceutical automation, GxP compliance, and validation. I help teams turn AI from a chatbot into a reliable operational layer.

[Book a meeting →](#)