




5 AI Video Models, Ranked by What an EU Studio Can Legally Ship in 2026

For an EU studio, ranking 2026's audio-native video models has almost nothing to do with clip length or price per second. It comes down to which licence lets you press render at all, and MiniMax H3 answers that with a no.

The short answer: LTX-2.5. Lightricks released open weights on 11 August 2026, free for commercial use under \$10M annual revenue, with no territorial carve-out reported. MiniMax H3 ranks last because its Community License, effective 2 August 2026, lists the EU as an Excluded Territory and Section V.4 extends that restriction to Outputs.

MiniMax published H3's open weights on Hugging Face on 3 August 2026, four days after the API launch. The Community License names the EU, UK, Republic of Korea and USA as Excluded Territories, and Section V.4 bars use, reproduction, modification, distribution or display of the model ****or its Outputs**** outside the Applicable Territory. The same day the licence took effect, EU AI Act Article 50 started applying. Those two dates are the whole story.



5 AI Video Models, Ranked by What an EU Studio Can Legally Ship in 2026

So the criterion for this ranking is narrow and deliberately unglamorous: what can a studio incorporated in the EU put into a client deliverable, with a paper trail, before 2 December 2026, when pre-existing generative systems must carry machine-readable marking or face fines up to €15M or 3% of worldwide annual turnover. Quality is not the axis. Six models shipped between 31 July and 27 August 2026, and the specs converged tightly enough that the technical spread stopped deciding anything.

ATTRIBUTE	LTX-2.5	WAN 3.0	FLUX 3 VIDEO	GEMINI OMNI 1.1 FLASH	MINIMAX H3
Weights available	Open, 11 Aug 2026	No published weights	API only	API only	Open, 3 Aug 2026
EU deployment allowed by licence	Yes, no territorial carve-out reported	API, no carve-out reported	API, no carve-out reported	API, no carve-out reported	No: EU is an Excluded Territory
Max single-pass clip	n/a in research	30s	20s	10s (40s chained)	15s
Entry price per second	\$0.09 (720p)	\$0.05 (480p)	\$0.12	~\$0.10 (720p)	\$0.08 (768p)
Top documented price per second	\$0.37 (4K, Pro)	\$0.20 (1080p)	\$0.53 (top tier)	n/a	\$0.13 (2K)
C2PA / default watermark documented	Not documented	n/a in research	n/a in research	n/a in research	Not documented

Two things about that table need saying out loud. “n/a in research” means exactly that: I have no sourced figure, and I am not going to estimate one for the sake of a full grid. And “no carve-out reported” is weaker than “permitted.” Absence of a reported restriction is not the same as a reviewed licence, and for API products the terms of service can change on a release note. If you are betting a production pipeline on one of these, read the current terms yourself rather than my summary of press coverage.



Can an EU studio legally run MiniMax H3 weights locally?

No, not under the Community License as published. The licence text breakdown from [Csidia's reading of Sections I.3, I.5 and V.4](#) puts the EU in the Excluded Territories list alongside the UK, South Korea and the USA, and MiniMax's own framing is that developers in excluded regions "must apply for a formal license." [TechTimes](#) and [AI Frontier Review](#) both reported the same carve-out within a day of the Hugging Face drop.

The part that should worry a studio more than the weights restriction is the Outputs clause. Section V.4 covers display and distribution of Outputs, not only of the model. Read plainly, that means a clip generated by someone in an Applicable Territory and then shown to a Munich client is still inside the restricted set. I am not a lawyer, and this is my reading of the reported text rather than legal advice. But it kills the obvious workaround, which was to have a partner outside the excluded zone render and hand over files.

There is a legitimate path: the API. H3's API launched 31 July 2026 at \$0.08/s for 768p and \$0.13/s for 2K, which is the cheapest 2K in this field by a wide margin. Whether the API terms carry the same territorial language as the weights licence is not something the research settles, and that gap is precisely where I would send a contract lawyer before spending money. I covered the model's technical envelope when it launched in [the H3 launch write-up](#): 33B parameters, 4 to 15 second clips, 768p base with 2K upscale, native stereo audio. Good model. Wrong passport.

The order I would actually deploy in

LTX-2.5 (Lightricks)

Open weights released 11 August 2026, free for commercial use under \$10M annual revenue, and no territorial carve-out reported: the only model here you can run on your own hardware, in the EU, and still own the pipeline.



Wan 3.0 (Alibaba)

API only, no published weights, but it does the thing nobody else does at that price: 2 to 30 second clips in one pass, up to 1080p, at \$0.05/\$0.10/\$0.20 per second for 480p/720p/1080p.

FLUX 3 Video (Black Forest Labs)

GA on 4 August 2026, 5 to 20 second clips with synchronized native audio, roughly \$0.12 to \$0.53 per second by tier. That is the most expensive top tier in the field, from a European vendor, which matters for contracting even though nothing in the research confirms provenance features.

Gemini Omni 1.1 Flash (Google)

Announced 27 August 2026 at around \$0.10/s for 720p, with 3 to 10 second generations chained in 10 second increments to a 40 second scene. The chaining is real capability, but Google retired the Imagen 4 generate endpoints on 17 August 2026, and that is a deprecation cadence you plan around.

MiniMax H3

Ranked last on one criterion only: the Community License effective 2 August 2026 names the EU an Excluded Territory, and Section V.4 extends that to Outputs.

LTX-2.5 at the top needs its caveats spelled out, because “run it yourself” is doing a lot of work in that verdict. The full model wants 80GB+ of VRAM. Distilled and FP8 variants run on 32GB, with a documented floor of 12GB, which is the number that makes local deployment



plausible for a small shop rather than a datacentre tenant. Lightricks reports a 10 second clip rendered in 6.8 seconds on two NVIDIA GB200 superchips, a vendor benchmark on hardware most studios will never touch, so treat it as a ceiling rather than a planning figure. If you prefer the hosted route, LTX API pricing runs \$0.09/s at 720p, \$0.15/s at 1080p, \$0.19/s at 2K and \$0.37/s at 4K on the Pro tier, up to 4K at 50fps. The \$10M revenue ceiling on free commercial use is generous for most studios and irrelevant for the ones it is not.

Wan 3.0's position is the one I am least comfortable with. It entered public beta on 6 August 2026 and went wide on 24 August, it has no published weights, and it is the cheapest per second in this field at 480p. That combination (API only, lowest price, a 30 second single pass) makes it excellent for volume work and a poor foundation for anything you need to reproduce in two years. Rank two is a verdict about throughput economics, not about sovereignty.

Provenance marking is your job, not the model's

Neither LTX-2.5 nor MiniMax H3 is documented as embedding C2PA Content Credentials or default watermarks. Article 50 has applied since 2 August 2026 and gives systems already on the market until 2 December 2026 for machine-readable marking, with exposure up to €15M or 3% of global turnover. If you run open weights, the marking obligation does not disappear. It moves to you and your post pipeline.

That is the trade the open-weights route actually demands. You gain a model nobody can deprecate out from under you, and you inherit a compliance workflow that the API vendors may or may not be handling on your behalf. The research does not confirm C2PA support for the API-only models either, so the safe assumption for all five is that provenance marking is your job until a vendor puts it in writing.

Seedance 2.5 misses the cut on an unpinnable resolution spec

ByteDance's Seedance 2.5 launched 31 July 2026 with a single continuous 30 second pass and native audio, which on paper is the most useful generation envelope in the field. It is off the ranking because the documented output in the BytePlus ModelArk docs is 480p and



720p, while higher claims conflict across hosts. A resolution spec you cannot pin down is not something I would put in a client SOW.

Worth watching rather than worth adopting. If ByteDance publishes a clear 1080p+ figure, it moves straight into contention with Wan 3.0 on the 30 second axis, and the choice between them becomes a licence question again rather than a capability one.

FLUX 3's video line deserves a note too. I wrote about [FLUX 3 at launch](#), and GA on 4 August confirmed the 20 second ceiling with synchronized audio. The top tier at \$0.53/s is more than ten times Wan 3.0's 480p rate. That spread is only defensible if the output quality gap justifies it, and the research here does not contain a head-to-head that would let me say.

My take: territorial licence clauses will spread

My take: six audio-native models in four weeks, and the differentiator turned out to be a clause in a licence file. Territorial licensing is going to spread, because it costs a vendor nothing to write and it lets them ship weights for prestige while keeping enforcement optional in the markets with real regulators. Expect more Excluded Territory lists, and expect at least one of them to arrive as a change to an existing licence rather than a new release. The defensive move is keeping your pipeline model-agnostic enough that a licence change costs you a week rather than a quarter. I would run LTX-2.5 weights locally as the fallback that cannot be revoked, and route production volume to whichever API is cheapest that month.


The 2 December 2026 deadline is the one date to put on a wall. Article 50 applied from 2 August, but systems already on the market get until December for machine-readable marking, per the [Ropes & Gray operational breakdown](#). That is roughly three months from the end of August to get provenance into an export step, confirm which of your vendors handles marking, and document the ones that do not. None of the five models in this ranking has documented C2PA support in the material I have, and I would not assume any of them will ship it for you.

WHAT TO DO WITH THIS

Pick on licence first, price second, clip length third. LTX-2.5 open weights under



the \$10M revenue ceiling is the only option here that an EU studio can run locally without a territorial question mark, and whatever you pick, the Article 50 marking work lands on 2 December 2026 regardless. If you want a second pair of eyes on which of these fits your contracts, [Book a call →](#)



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 →