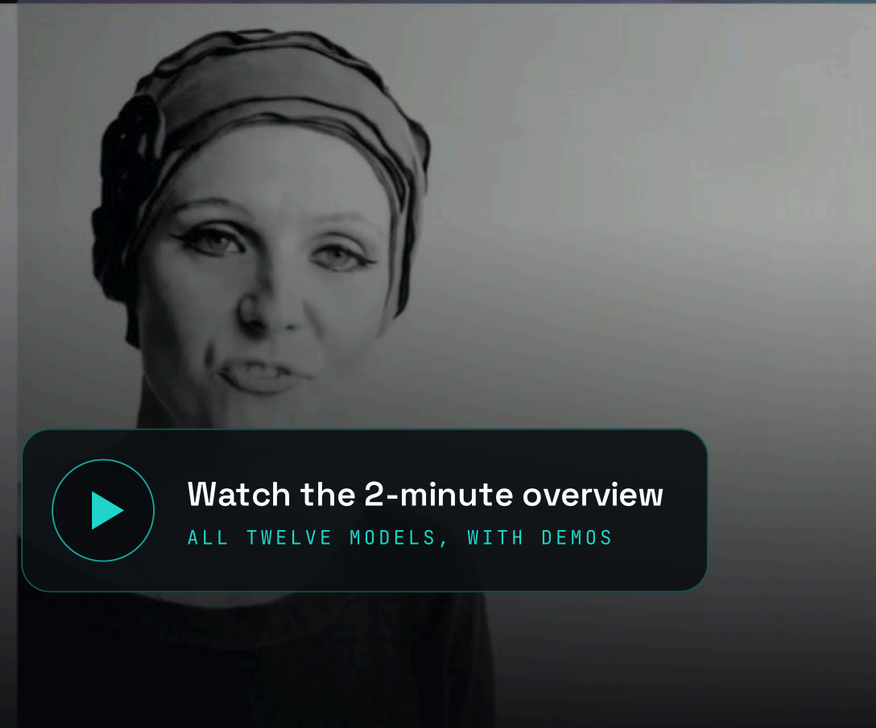
iK INFERENCEKEY

Open video models your brand can actually use

Twelve open video models a large EU company can run on its own GPUs and use commercially, what each one costs per hour, and a 90-day pilot to prove it on your own footage.

PREPARED FOR	DATE	FROM
ADIDAS	NOVEMBER 2026	INFERENCEKEY





Watch the 2-minute overview

ALL TWELVE MODELS, WITH DEMOS

01 CLEARED FOR A EUROPEAN COMPANY

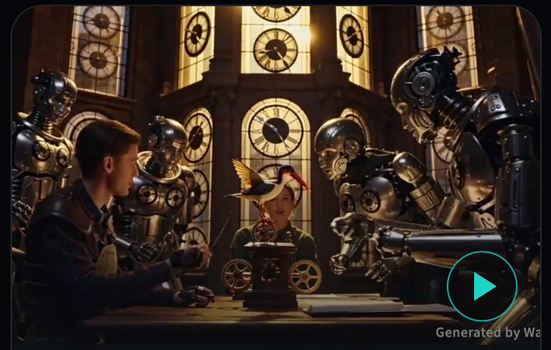
These are the open video models a European company can use

Apache-2.0 or MIT: commercial use, no revenue cap, no territory clause. What each one does, the memory it needs and what the card costs.

MODEL	WHAT IT DOES	LICENCE	GPU MEMORY	CHEAPEST CARD · $/H	COMFORTABLE CARD · $/H
Wan 2.2 A14B	Text or image to video	Apache-2.0	41 GB at 480p, 60 GB at 720p	A40 48 GB · **0.59**	A100 80 GB · **2.09**
Wan 2.2 TI2V-5B	Fast 720p drafts	Apache-2.0	23 GB	RTX A5000 24 GB · **0.27**	RTX 4090 24 GB · **0.89**
LongCat-Video	Long shots, up to minutes	MIT	about 48 GB	A40 48 GB · **0.59**	H100 80 GB · **2.89**
Ovi	Video with speech and sound	Apache-2.0	32 GB (24 GB in FP8)	RTX PRO 4500 32 GB · **0.72**	H100 80 GB · **2.89**
daVinci-MagiHuman	People talking on camera, 1080p	Apache-2.0	80 GB	H100 80 GB · **2.89**	H100 SXM 80 GB · **4.39**
JoyAI-Video-Edit 1.1 FASHION PICK	Swap garments and products in real footage	Apache-2.0	28 GB at 480p, 47 GB at 720p	RTX 5090 32 GB · **1.39**	RTX PRO 6000 96 GB · **2.69**
Wan-Animate-2	Animate a mascot from a dancer's video	Apache-2.0	about 24 GB	RTX A5000 24 GB · **0.27**	A100 80 GB · **2.09**
SCAIL-2	Several characters, or replace one person	MIT	about 24 GB	RTX A5000 24 GB · **0.27**	A100 80 GB · **2.09**
Wan2.1 VACE	Fill, remove or extend parts of a shot	Apache-2.0	about 24 GB	RTX A5000 24 GB · **0.27**	A100 80 GB · **2.09**
InfiniteTalk	Re-voice existing footage in another language	Apache-2.0	about 24 GB	RTX A5000 24 GB · **0.27**	A100 80 GB · **2.09**
LongCat-Video-Avatar 1.5	Spokesperson from a photo and a voice	MIT	about 24 GB	RTX A5000 24 GB · **0.27**	2x A100 80 GB · **4.18**
FlashVSR+	Upscale to 4K	Apache-2.0	16 to 26 GB	RTX A5000 24 GB · **0.27**	RTX 4090 24 GB · **0.89**

On-demand data-centre prices on our cloud, US dollars per GPU per hour, checked on 13 November 2026. "Cheapest" runs the model, usually slowly. "Comfortable" runs it without fighting for room. "About" means the vendor publishes no official figure. Only two are tied to a chip family: daVinci-MagiHuman to Hopper (H100), and JoyAI-Video-Edit's real-time setups to Blackwell (RTX 5090, RTX PRO 6000, B200).

A clip from a sentence



ALIBABA · WAN TEAM

Wan 2.2

The safe pick, with the biggest ecosystem of tools. One 5-second 720p clip takes about 17 minutes on one H100. Its small sibling, TI2V-5B, fits a 24 GB card.

Apache-2.0 · 41 to 60 GB · A40 $0.59/h



MEITUAN

LongCat-Video

Keeps extending one shot for minutes. Scores level with Wan 2.2 in Meituan's own tests. 142 seconds for a 720p clip on one H800.

MIT · about 48 GB · A40 $0.59/h



CHARACTER.AI · YALE

Ovi

Writes picture and sound in one pass, dialogue included.

Apache-2.0 · 32 GB · RTX PRO 4500 $0.72/h

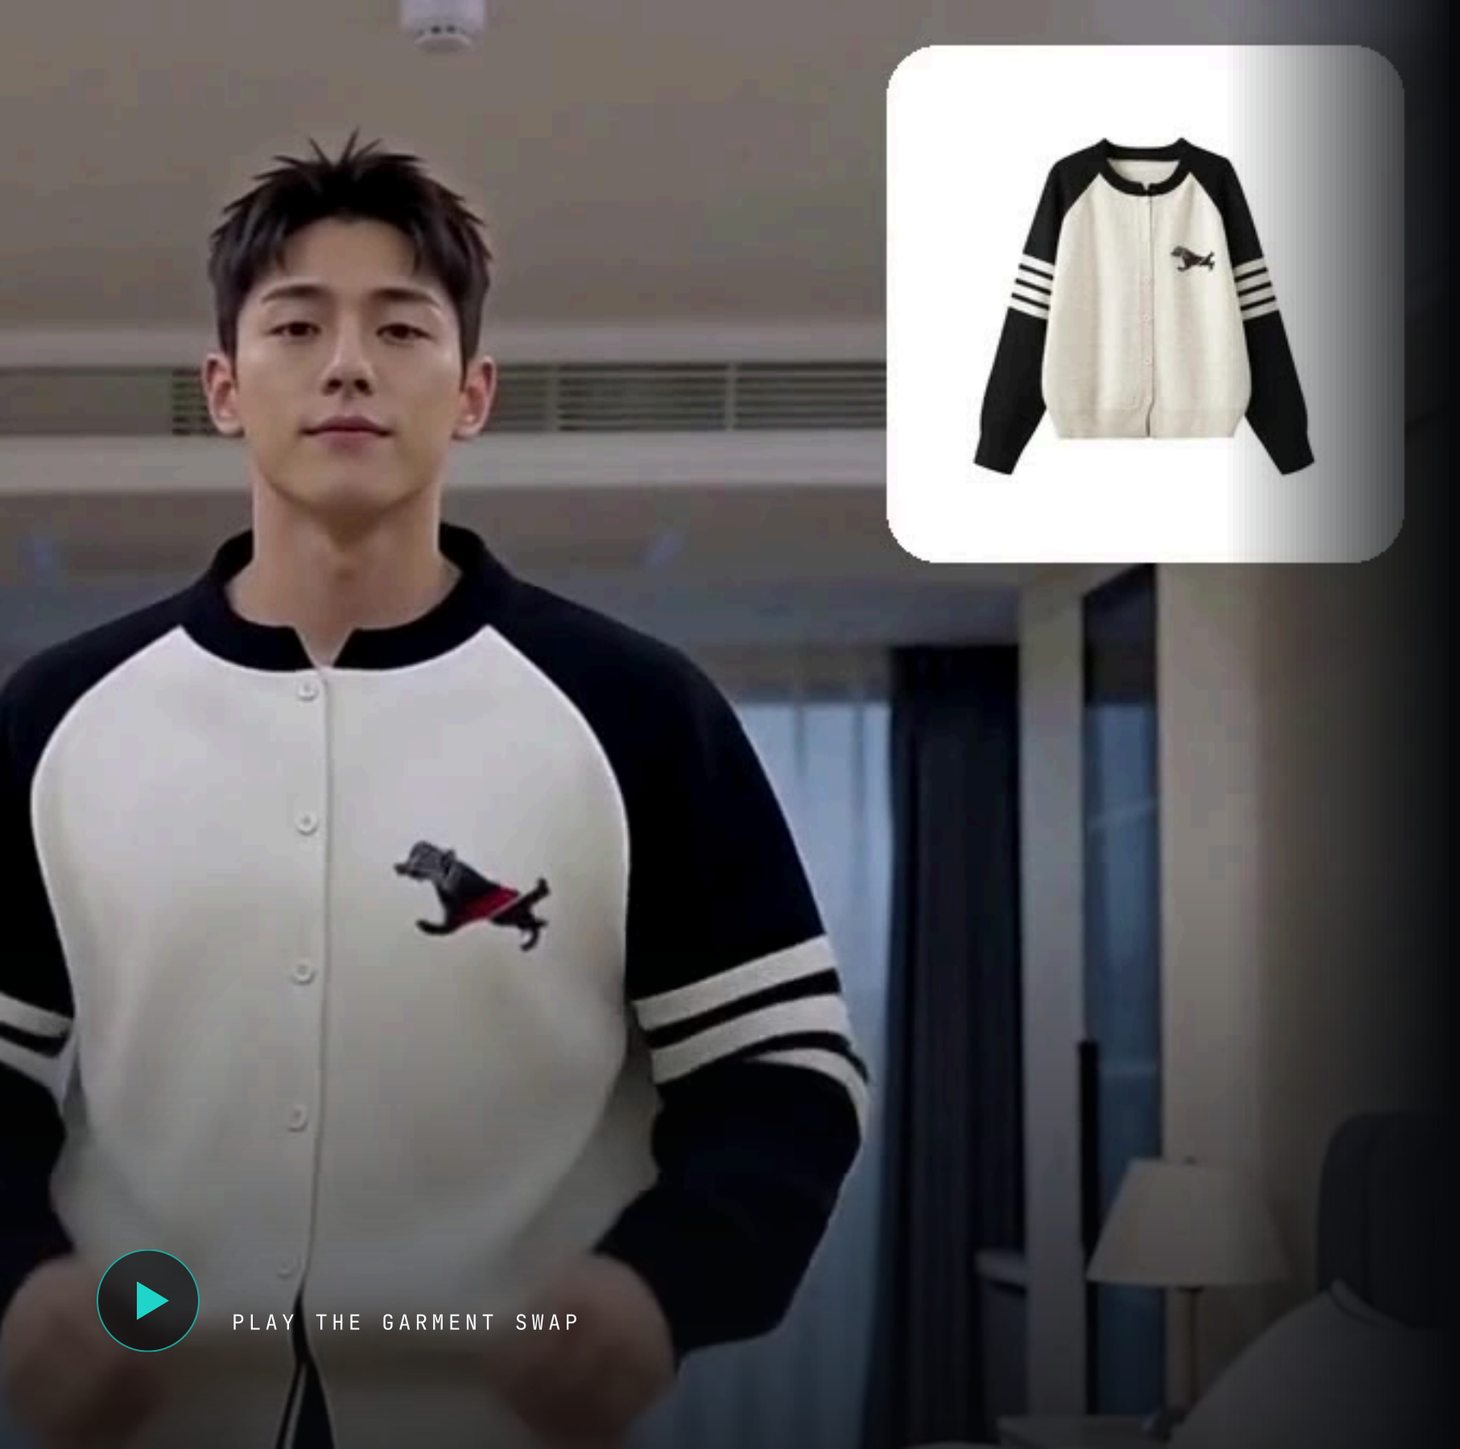


SII-GAIR · SAND.AI

daVinci-MagiHuman

The fast one for people on camera: 5 seconds of 1080p with speech in 38 seconds on an H100, about 3 cents a clip.

Apache-2.0 · 80 GB · Hopper · H100 $2.89/h



03 THE FASHION USE CASE

A product photo in. The model wears it.

01 Take a product photo of the garment and an existing video of a model.

02 JoyAI-Video-Edit dresses the model in that garment while the movement stays.

03 It also takes plain instructions: "dress the girl in a brown down jacket and a blue cap".



Apache-2.0 · JD.com

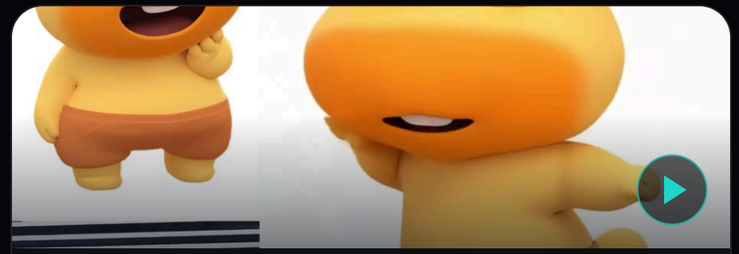
28 GB at 480p, 47 GB at 720p

RTX 5090 $1.39/h

RTX PRO 6000 $2.69/h

Real time on Blackwell

Work on footage you already have



Wan-Animate-2

Makes a mascot dance: one picture of the character, one video of a person.

Apache-2.0 | about 24 GB | RTX A5000 $0.27/h



SCAIL-2

The same with several characters, or swaps one person for another.

MIT | about 24 GB | RTX A5000 $0.27/h



Wan2.1 VACE

Fill a masked area, remove an object, or turn a 16:9 ad into a 9:16 story.

Apache-2.0 | about 24 GB | RTX A5000 $0.27/h



InfiniteTalk

Re-voices existing footage: new audio in, lips, head and body follow. One shoot, many languages.

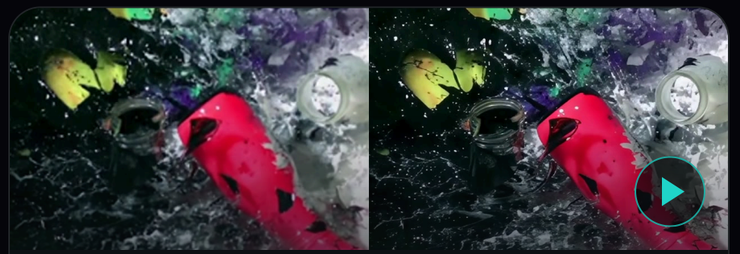
Apache-2.0 | about 24 GB | RTX A5000 $0.27/h



LongCat-Video-Avatar 1.5

A spokesperson who talks, from one photo and one voice.

MIT | about 24 GB | RTX A5000 $0.27/h



FlashVSR+

Takes a clip to 4K on one 24 GB card, about one frame per second on an RTX 4090.

Apache-2.0 | 16 to 26 GB | RTX A5000 $0.27/h

One platform, one endpoint, every video model.

Your teams
Marketing, e-commerce, content studios

→

INFERENCEKEY
One endpoint, with projects, keys and quotas per team

→

- The open video models — commercial use
- **Dedicated GPUs** — IN OUR CLOUD
- Your own workflows — e.g. garment swap

SIMPLE TO ADOPT

Call it like any API

Your teams send a request and get the clip back. Each team has its own project, key and quota, and every job is metered.

PAY FOR WHAT RUNS

Scales to zero when idle

A GPU starts when a job arrives and stops when the work is done. No machine runs, and nothing is billed, while nobody uses it.

Try it on your own footage

We set it up and run it. Your team only has to use it, and nothing changes in the tools you already have.

FOR YOUR TEAMS

Nothing to change

- Works alongside the tools you already use
- One place to reach every model
- Start with one use case, such as garment swaps

WE DO

We run it for you

- A private space for adidas, managed by us
- The models you pick, ready to use
- A direct line to our engineers

HOW YOU START

Your choice

- Month to month, or a yearly plan
- On the yearly plan, setup is on us
- GPU billed only while it is in use

See the models in motion



The 2-minute overview

All twelve models, with demos



The garment swap

JoyAI-Video-Edit

iK inferencekey.com

Licences: this is our reading of each model's public licence as of November 2026, not legal advice. Run the final pick past your legal team. The demo clips come from each vendor's repo or project page, carry no licence of their own and stay with their authors; treat them as examples of what each model does. GPU prices are on-demand rates on our cloud on 13 November 2026 and can change.