GPUpartner

VALIDATION REPORT · VIDEO GENERATION

# Wan2.1 I2V 14B

on Radeon RX 7900 XTX · 8-bit (FP8 · Q8_0)

STATUS · NOV 22, 2026

**Not validated**

2 of 3 checks pass

**Wan2.1 I2V 14B at 8-bit runs on the Radeon RX 7900 XTX, about 71 min per clip.**

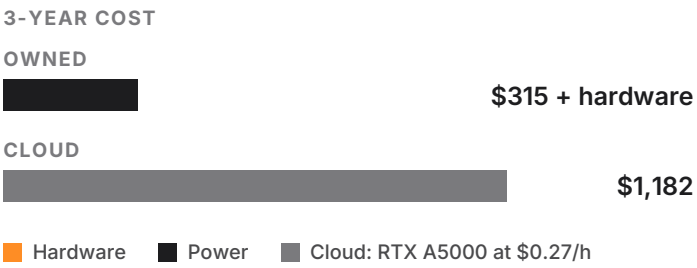
## VALIDATION

| CHECK | TARGET | RESULT | STATUS |
|---|---|---|---|
| Fits in GPU memory | 24.0 GB usable | 20.4 GB needed | PASS |
| Headroom | 5% free | 3.6 GB free (15%) | PASS |
| Time per clip | under 15 min | 71 min | FAIL |

Targets: a clip in under 15 min, and 5% of memory left free. Calculated from published specs, not benchmarked.

## BUDGET

| | |
|---|---|
| HARDWARE, ONE-TIME | MSRP not entered |
| POWER, PER YEAR | $105 · 400 W at the wall |
| CLOUD INSTEAD, PER YEAR | $394 |
| 3-YEAR TOTAL, OWNED | Hardware + $315 power |
| 3-YEAR TOTAL, CLOUD | $1,182 |
| BREAK-EVEN | Needs the hardware MSRP |




Budget basis: 4 h a day, power at $0.18 per kWh, cloud at RunPod on-demand prices as of Nov 22, 2026. Hardware at the MSRP entered in the calculator, not a GPU Partner quote. A 3-year view; resale value, staff time and taxes aren't counted.

GPUpartner

TECHNICAL DETAIL · VIDEO GENERATION

# Wan2.1 I2V 14B

on Radeon RX 7900 XTX · 8-bit (FP8 · Q8_0)

## SETUP

| | |
|---|---|
| PARAMETERS | 14.3B |
| PRECISION | 8-bit (FP8 · Q8_0) |
| OUTPUT | 1280 × 720, 81 frames |
| STEPS | 40 |
| HARDWARE | Radeon RX 7900 XTX |
| GPU CAN USE | 24.0 GB of 24.0 GB |

## MEMORY

| | |
|---|---|
| WEIGHTS | 13.8 GB |
| WORKING | 5.8 GB |
| RUNTIME | 819 MB |
| SYSTEM | 0.0 GB |
| FREE | 3.6 GB |

## SAME MODEL, OTHER MACHINES

| | |
|---|---|
| RTX 4090 | 8-bit · 31 min |
| RTX 5090 | 8-bit · 24 min |
| RTX PRO 6000 BLACKWELL | 16-bit · 10 min |
| M5 MAX 128 GB | 16-bit · 1.6 h |
| M3 ULTRA 256 GB | 16-bit · 3.1 h |
| DGX SPARK 128 GB | 16-bit · 81 min |
| RYZEN AI MAX+ 395 128 GB | 16-bit · 2.4 h |

## WHERE TO GET IT

| | |
|---|---|
| OFFICIAL | huggingface.co/Wan-AI/Wan2.1-I2V-14B-720P-Diffusers |

## METHOD AND ASSUMPTIONS

Memory is the denoiser's weights at the chosen precision plus the bigger of one transformer block's working set and the VAE decode, plus the runtime; the text encoders run first and then wait in system RAM. Time is the denoising FLOPs per step times the steps, against the chip's dense 16-bit TFLOPS, scaled to published GeForce runs of this model. Running cost is power at the wall at your rate and hours; the cloud figure is the cheapest RunPod on-demand GPU that runs the same job. Accuracy: memory lands within about a gigabyte; speeds and times are rough, about ±25%. Model data and hardware specs checked Nov 2026; cloud prices are RunPod's on-demand list as of Nov 22, 2026. Report VR-1V1VU8V, generated November 22, 2026.