GPUpartner

VALIDATION REPORT · VIDEO GENERATION

MiniMax H3

on Radeon RX 7900 XTX · 16-bit (BF16)

STATUS · NOV 22, 2026

Not validated

0 of 2 checks pass

MiniMax H3 at 16-bit needs 40.9 GB; the Radeon RX 7900 XTX can use 24.0 GB. The smallest machine we checked that runs it is the RTX PRO 5000 Blackwell 48 GB.

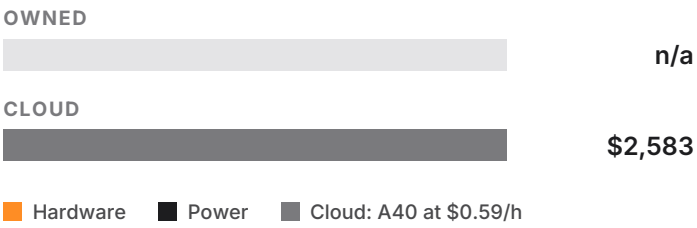
VALIDATION

CHECK	TARGET	RESULT	STATUS
Fits in GPU memory	24.0 GB usable	40.9 GB needed	FAIL
Headroom	5% free	none	FAIL
Time per clip	under 15 min	–	NOT TESTED

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	Doesn't run on this machine
CLOUD INSTEAD, PER YEAR	$861
3-YEAR TOTAL, OWNED	n/a
3-YEAR TOTAL, CLOUD	$2,583
BREAK-EVEN	n/a



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

MiniMax H3

on Radeon RX 7900 XTX · 16-bit (BF16)

SETUP

PARAMETERS	20.1B
PRECISION	16-bit (BF16)
HARDWARE	Radeon RX 7900 XTX
GPU CAN USE	24.0 GB of 24.0 GB

MEMORY

WEIGHTS	37.5 GB
WORKING	2.6 GB
RUNTIME	819 MB
SYSTEM	0.0 GB
OVER BY	16.9 GB

SUGGESTED CHANGES

SPILL TO RAM	Runs, slower, 14 min a clip
RTX PRO 5000 BLACKWELL 48 GB	Runs it, 4.0 min a clip
M5 MAX · 128 GB	Runs it, 18 min a clip
RYZEN AI MAX+ 395 · 128 GB	Runs it, 28 min a clip

SAME MODEL, OTHER MACHINES

RTX 4090	Too big
RTX 5090	8-bit · 4.9 min
RTX PRO 6000 BLACKWELL	16-bit · 2.2 min
M5 MAX 128 GB	16-bit · 18 min
M3 ULTRA 256 GB	16-bit · 35 min
DGX SPARK 128 GB	16-bit · 16 min
RYZEN AI MAX+ 395 128 GB	16-bit · 28 min

WHERE TO GET IT

FP8	huggingface.co/Comfy-Org/MiniMax-H3
NVFP4	huggingface.co/lilcheaty/MiniMax-H3-NVFP4
GGUF	huggingface.co/Abiray/MiniMax-H3-GGUF
OFFICIAL	huggingface.co/MiniMaxAI/MiniMax-H3

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-18GU9M9, generated November 22, 2026.