GPUpartner

VALIDATION REPORT · IMAGE GENERATION

FLUX.2 [dev]

on Radeon AI PRO R9700 · 6-bit (GGUF Q6_K)

STATUS · NOV 22, 2026

Not validated

2 of 3 checks pass

FLUX.2 [dev] at 6-bit runs on the Radeon AI PRO R9700, about 2.6 min per image.

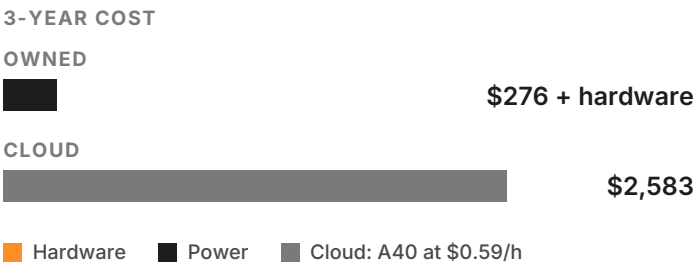
VALIDATION

CHECK	TARGET	RESULT	STATUS
Fits in GPU memory	32.0 GB usable	26.0 GB needed	PASS
Headroom	5% free	6.0 GB free (19%)	PASS
Time per image	under 30 s	2.6 min	FAIL

Targets: an image in under 30 s, and 5% of memory left free. Calculated from published specs, not benchmarked.

BUDGET

HARDWARE, ONE-TIME	MSRP not entered
POWER, PER YEAR	$92 · 350 W at the wall
CLOUD INSTEAD, PER YEAR	$861
3-YEAR TOTAL, OWNED	Hardware + $276 power
3-YEAR TOTAL, CLOUD	$2,583
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 · IMAGE GENERATION

FLUX.2 [dev]

on Radeon AI PRO R9700 · 6-bit (GGUF Q6_K)

SETUP

PARAMETERS	32.2B
PRECISION	6-bit (GGUF Q6_K)
OUTPUT	1024 × 1024
STEPS	28
HARDWARE	Radeon AI PRO R9700
GPU CAN USE	32.0 GB of 32.0 GB

MEMORY

WEIGHTS	24.8 GB
WORKING	432 MB
RUNTIME	819 MB
SYSTEM	0.0 GB
FREE	6.0 GB

SAME MODEL, OTHER MACHINES

RTX 4090	5-bit · 1.8 min
RTX 5090	6-bit · 85 s
RTX PRO 6000 BLACKWELL	16-bit · 31 s
M5 MAX 128 GB	16-bit · 4.8 min
M3 ULTRA 256 GB	16-bit · 9.3 min
DGX SPARK 128 GB	16-bit · 4.1 min
RYZEN AI MAX+ 395 128 GB	16-bit · 7.3 min

WHERE TO GET IT

FP8	huggingface.co/Comfy-Org/flux2-dev
NVFP4	huggingface.co/black-forest-labs/FLUX.2-dev-NVFP4
GGUF	huggingface.co/city96/FLUX.2-dev-gguf
OFFICIAL	huggingface.co/black-forest-labs/FLUX.2-dev

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-0BK7AVM, generated November 22, 2026.