

VALIDATION REPORT · IMAGE GENERATION

FLUX.2 [dev]

on RTX 4060 Ti 16 GB · 16-bit (BF16)

STATUS · OCT 22, 2026

Not validated

0 of 2 checks pass

FLUX.2 [dev] at 16-bit needs 61.2 GB; the RTX 4060 Ti 16 GB can use 16.0 GB. The smallest machine we checked that runs it is the RTX PRO 5000 Blackwell 72 GB.

VALIDATION

CHECK	TARGET	RESULT	STATUS
Fits in GPU memory	16.0 GB usable	61.2 GB needed	FAIL
Headroom	5% free	none	FAIL
Time per image	under 30 s	–	NOT TESTED


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	Doesn't run on this machine
CLOUD INSTEAD, PER YEAR	\$2,612
3-YEAR TOTAL, OWNED	n/a
3-YEAR TOTAL, CLOUD	\$7,836
BREAK-EVEN	n/a

3-YEAR COST

OWNED

 n/a

CLOUD

 \$7,836

Hardware Power Cloud: A100 80GB at \$1.79/h

Budget basis: 4 h a day, power at \$0.18 per kWh, cloud at RunPod on-demand prices as of Oct 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.

TECHNICAL DETAIL · IMAGE GENERATION

FLUX.2 [dev]

on RTX 4060 Ti 16 GB · 16-bit (BF16)

SETUP

PARAMETERS	32.2B
PRECISION	16-bit (BF16)
HARDWARE	RTX 4060 Ti 16 GB
GPU CAN USE	16.0 GB of 16.0 GB

MEMORY

WEIGHTS	60.0 GB
WORKING	432 MB
RUNTIME	819 MB
SYSTEM	0.0 GB
OVER BY	45.2 GB

SUGGESTED CHANGES

SPILL TO RAM	Runs, slower, 6.7 min an image
RTX PRO 5000 BLACKWELL 72 GB	Runs it, 60 s an image
M5 MAX · 128 GB	Runs it, 4.8 min an image
RYZEN AI MAX+ 395 · 128 GB	Runs it, 7.3 min an image

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 $\pm 25\%$. Model data and hardware specs checked Oct 2026; cloud prices are RunPod's on-demand list as of Oct 22, 2026. Report VR-09QX7L7, generated October 22, 2026.