

VALIDATION REPORT · IMAGE GENERATION

FLUX.1 [dev]

on RTX 4090 · 16-bit (BF16)

STATUS · OCT 22, 2026

Validated with conditions

2 of 3 checks pass

FLUX.1 [dev] at 16-bit runs on the RTX 4090, about 12 s per image.

VALIDATION

CHECK	TARGET	RESULT	STATUS
Fits in GPU memory	24.0 GB usable	23.2 GB needed	PASS
Headroom	5% free	843 MB free (3%)	CAUTION
Time per image	under 30 s	12 s	PASS

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	\$127 · 485 W at the wall
CLOUD INSTEAD, PER YEAR	\$394
3-YEAR TOTAL, OWNED	Hardware + \$382 power
3-YEAR TOTAL, CLOUD	\$1,182
BREAK-EVEN	Needs the hardware MSRP

3-YEAR COST

OWNED



\$382 + hardware

CLOUD



\$1,182

Hardware Power Cloud: RTX A5000 at \$0.27/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.1 [dev]

on RTX 4090 · 16-bit (BF16)

SETUP

PARAMETERS	11.9B
PRECISION	16-bit (BF16)
OUTPUT	1024 × 1024
STEPS	20
HARDWARE	RTX 4090
GPU CAN USE	24.0 GB of 24.0 GB

MEMORY

WEIGHTS	22.2 GB
WORKING	216 MB
RUNTIME	819 MB
SYSTEM	0.0 GB
FREE	843 MB

SUGGESTED CHANGES

8-BIT WEIGHTS	Fits, 12 s an image
768 ² OUTPUT	Fits, 7.3 s an image
RTX 5090	Runs it, 9.5 s an image
M5 MAX · 48 GB	Runs it, 35 s an image

SAME MODEL, OTHER MACHINES

RTX 5090	16-bit · 9.5 s
RTX PRO 6000 BLACKWELL	16-bit · 4.6 s
M5 MAX 128 GB	16-bit · 35 s
M3 ULTRA 256 GB	16-bit · 66 s
DGX SPARK 128 GB	16-bit · 30 s
RYZEN AI MAX+ 395 128 GB	16-bit · 53 s
DGX STATION GB300	16-bit · 1.7 s

WHERE TO GET IT

FP8	huggingface.co/black-forest-labs/FLUX.1-dev-FP8
NVFP4	huggingface.co/black-forest-labs/FLUX.1-dev-NVFP4
GGUF	huggingface.co/city96/FLUX.1-dev-gguf
OFFICIAL	huggingface.co/black-forest-labs/FLUX.1-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 Oct 2026; cloud prices are RunPod's on-demand list as of Oct 22, 2026. Report VR-1V7MTEG, generated October 22, 2026.