

VALIDATION REPORT · IMAGE GENERATION

FLUX.2 [klein] 9B

on Ryzen AI Max+ 395 128 GB · 16-bit (BF16)

STATUS · OCT 22, 2026

Validated

3 of 3 checks pass

FLUX.2 [klein] 9B at 16-bit runs on the Ryzen AI Max+ 395 128 GB, about 13 s per image.

VALIDATION

CHECK	TARGET	RESULT	STATUS
Fits in GPU memory	96.0 GB usable	27.7 GB needed	PASS
Headroom	5% free	68.3 GB free (71%)	PASS
Time per image	under 30 s	13 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	\$33 · 126 W at the wall
CLOUD INSTEAD, PER YEAR	\$365
3-YEAR TOTAL, OWNED	Hardware + \$99 power
3-YEAR TOTAL, CLOUD	\$1,094
BREAK-EVEN	Needs the hardware MSRP

3-YEAR COST

OWNED



\$99 + hardware

CLOUD



\$1,094

■ Hardware ■ Power ■ Cloud: RTX A4500 at \$0.25/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 [klein] 9B

on Ryzen AI Max+ 395 128 GB · 16-bit (BF16)

SETUP

PARAMETERS	9.1B
PRECISION	16-bit (BF16)
OUTPUT	1024 × 1024
STEPS	4
HARDWARE	Ryzen AI Max+ 395 128 GB
GPU CAN USE	96.0 GB of 128 GB

MEMORY

WEIGHTS	24.5 GB
WORKING	2.3 GB
RUNTIME	819 MB
SYSTEM	32.0 GB
FREE	68.3 GB

SAME MODEL, OTHER MACHINES

RTX 4090	16-bit · 3.5 s
RTX 5090	16-bit · 3.0 s
RTX PRO 6000 BLACKWELL	16-bit · 1.8 s
M5 MAX 128 GB	16-bit · 8.8 s
M3 ULTRA 256 GB	16-bit · 16 s
DGX SPARK 128 GB	16-bit · 7.6 s
DGX STATION GB300	16-bit · 1.2 s

WHERE TO GET IT

FP8	huggingface.co/black-forest-labs/FLUX.2-klein-9b-fp8
NVFP4	huggingface.co/black-forest-labs/FLUX.2-klein-9b-nvfp4
GGUF	huggingface.co/unsloth/FLUX.2-klein-9B-GGUF
OFFICIAL	huggingface.co/black-forest-labs/FLUX.2-klein-9B

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-1CBQGM3, generated October 22, 2026.