GPUpartner

VALIDATION REPORT · TEXT GENERATION

Kimi VL A3B Thinking 2506

on RTX 4060 · 4-bit (Q4_K_M) · Ollama / LM Studio · Chat

STATUS · NOV 22, 2026

Not validated

0 of 2 checks pass

Kimi VL A3B Thinking 2506 at 4-bit needs 10.3 GB; the RTX 4060 can use 8.0 GB. The smallest machine we checked that runs it is the RTX PRO 2000 Blackwell.

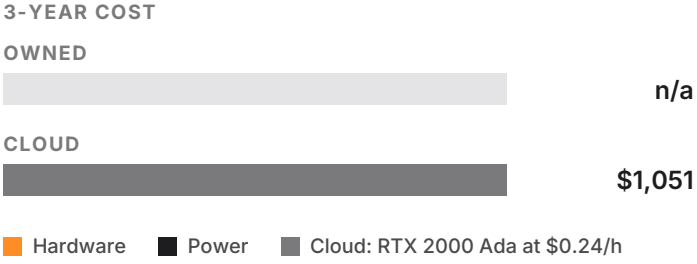
VALIDATION

CHECK	TARGET	RESULT	STATUS
Fits in GPU memory	8.0 GB usable	10.3 GB needed	FAIL
Headroom	5% free	none	FAIL
Context for chat	16K or more	–	NOT TESTED
Writing speed	15 tok/s or more	–	NOT TESTED
First reply	under 3 s, 2K prompt	–	NOT TESTED

Targets: a reply you can read along with, a first reply the use case tolerates, 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	$350
3-YEAR TOTAL, OWNED	n/a
3-YEAR TOTAL, CLOUD	$1,051
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 · TEXT GENERATION

Kimi VL A3B Thinking 2506

on RTX 4060 · 4-bit (Q4_K_M) · Ollama / LM Studio · Chat

SETUP

PARAMETERS	16.4B, 3B active
PRECISION	4-bit (Q4_K_M)
CONTEXT	8K, auto
CHATS AT ONCE	1
USE	Chat
RUNS IN	Ollama / LM Studio
HARDWARE	RTX 4060
GPU CAN USE	8.0 GB of 8.0 GB

MEMORY

WEIGHTS	9.3 GB
CONTEXT	243 MB
RUNTIME	845 MB
SYSTEM	0.0 GB
OVER BY	2.3 GB

SUGGESTED CHANGES

2× RTX 4060	Fits across the cards, 128K context
SPILL TO RAM	Runs, slower, 48 tok/s
RTX PRO 2000 BLACKWELL	Runs it with 128K context
M5 MAX · 36 GB	Runs it with 128K context

SAME MODEL, OTHER MACHINES

RTX 4090	8-bit · 163 tok/s
RTX 5090	8-bit · 240 tok/s
RTX PRO 6000 BLACKWELL	8-bit · 240 tok/s
M5 MAX 128 GB	8-bit · 93 tok/s
M3 ULTRA 256 GB	8-bit · 102 tok/s
DGX SPARK 128 GB	8-bit · 55 tok/s
RYZEN AI MAX+ 395 128 GB	8-bit · 43 tok/s

HOW TO RUN IT

LLAMA.CPP	`llama-server -hf ggml-org/Kimi-VL-A3B-Thinking-2506-GGUF`
VLLM	`vllm serve moonshotai/Kimi-VL-A3B-Thinking-2506`
MODEL CARD	huggingface.co/moonshotai/Kimi-VL-A3B-Thinking-2506

METHOD AND ASSUMPTIONS

Memory is the weights (parameters times bits per weight, read from each model's config.json), the context cache for every chat, and the runtime, against the memory this GPU can use after the OS and display take theirs. Writing speed is memory bandwidth over the bytes read per token, plus a per-layer latency, calibrated against llama.cpp benchmarks. First reply is the prompt's FLOPs over the chip's dense 16-bit TFLOPS at the runtime's measured efficiency. 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-046PUDB, generated November 22, 2026.