GPUpartner

VALIDATION REPORT · TEXT GENERATION

Qwen3 VL 4B Thinking

on RTX 4090 · 8-bit (Q8_0) · Ollama / LM Studio · Chat

STATUS · NOV 22, 2026

Validated

5 of 5 checks pass

Qwen3 VL 4B Thinking at 8-bit runs on the RTX 4090 with 120K of context at about 129 tokens a second.

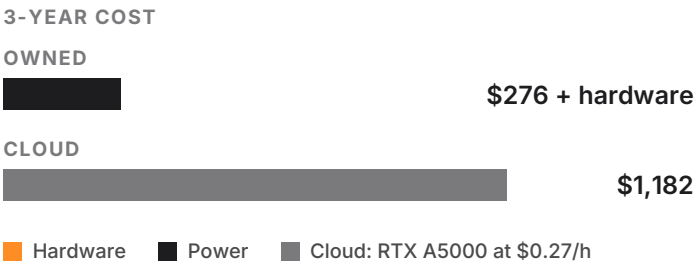
VALIDATION

CHECK	TARGET	RESULT	STATUS
Fits in GPU memory	24.0 GB usable	**22.1 GB needed**	PASS
Headroom	5% free	**1.9 GB free (8%)**	PASS
Context for chat	16K or more	**120K**	PASS
Writing speed	15 tok/s or more	**129 tok/s**	PASS
First reply	under 3 s, 2K prompt	**0.1 s**	PASS

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	$92 · 350 W at the wall
CLOUD INSTEAD, PER YEAR	$394
3-YEAR TOTAL, OWNED	Hardware + $276 power
3-YEAR TOTAL, CLOUD	$1,182
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 · TEXT GENERATION

Qwen3 VL 4B Thinking

on RTX 4090 · 8-bit (Q8_0) · Ollama / LM Studio · Chat

SETUP

PARAMETERS	4.4B
PRECISION	8-bit (Q8_0)
CONTEXT	120K, auto
CHATS AT ONCE	1
USE	Chat
RUNS IN	Ollama / LM Studio
HARDWARE	RTX 4090
GPU CAN USE	24.0 GB of 24.0 GB

SAME MODEL, OTHER MACHINES

RTX 5090	8-bit · 207 tok/s
RTX PRO 6000 BLACKWELL	8-bit · 207 tok/s
M5 MAX 128 GB	8-bit · 76 tok/s
M3 ULTRA 256 GB	8-bit · 85 tok/s
DGX SPARK 128 GB	8-bit · 39 tok/s
RYZEN AI MAX+ 395 128 GB	8-bit · 31 tok/s
DGX STATION GB300	8-bit · 493 tok/s

MEMORY

WEIGHTS	4.4 GB
CONTEXT	16.9 GB
RUNTIME	838 MB
SYSTEM	0.0 GB
FREE	1.9 GB

HOW TO RUN IT

LLAMA.CPP	`llama-server -hf unsloth/Qwen3-VL-4B-Thinking-GGUF`
VLLM	`vllm serve Qwen/Qwen3-VL-4B-Thinking`
MODEL CARD	huggingface.co/Qwen/Qwen3-VL-4B-Thinking

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-1GHEUG7, generated November 22, 2026.