GPUpartner

VALIDATION REPORT · TEXT GENERATION

# Qwen3 VL 8B Thinking

on RTX 4060 · 3-bit (Q3_K_M) · Ollama / LM Studio · Chat

STATUS · NOV 22, 2026

**Validated**

5 of 5 checks pass

**Qwen3 VL 8B Thinking at 3-bit runs on the RTX 4060 with 16K of context at about 45 tokens a second.**

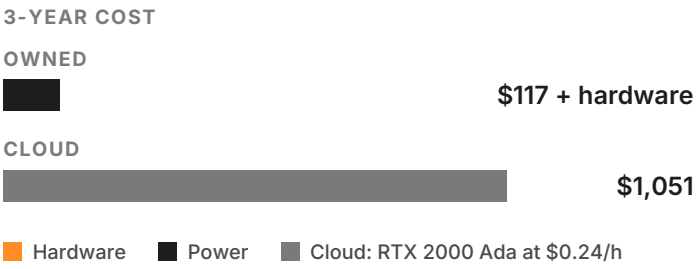
## VALIDATION

| CHECK | TARGET | RESULT | STATUS |
|---|---|---|---|
| Fits in GPU memory | 8.0 GB usable | 7.1 GB needed | PASS |
| Headroom | 5% free | 916 MB free (11%) | PASS |
| Context for chat | 16K or more | 16K | PASS |
| Writing speed | 15 tok/s or more | 45 tok/s | PASS |
| First reply | under 3 s, 2K prompt | 1.1 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 | $39 · 149 W at the wall |
| CLOUD INSTEAD, PER YEAR | $350 |
| 3-YEAR TOTAL, OWNED | Hardware + $117 power |
| 3-YEAR TOTAL, CLOUD | $1,051 |
| 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 8B Thinking

on RTX 4060 · 3-bit (Q3_K_M) · Ollama / LM Studio · Chat

## SETUP

| | |
|---|---|
| PARAMETERS | 8.8B |
| PRECISION | 3-bit (Q3_K_M) |
| CONTEXT | 16K, 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 | 4.0 GB |
| CONTEXT | 2.3 GB |
| RUNTIME | 886 MB |
| SYSTEM | 0.0 GB |
| FREE | 916 MB |

## SUGGESTED CHANGES

| | |
|---|---|
| 8-BIT CACHE | Room for 36K context |

## SAME MODEL, OTHER MACHINES

| | |
|---|---|
| RTX 4090 | 8-bit · 75 tok/s |
| RTX 5090 | 8-bit · 125 tok/s |
| RTX PRO 6000 BLACKWELL | 8-bit · 125 tok/s |
| M5 MAX 128 GB | 8-bit · 46 tok/s |
| M3 ULTRA 256 GB | 8-bit · 52 tok/s |
| DGX SPARK 128 GB | 8-bit · 22 tok/s |
| RYZEN AI MAX+ 395 128 GB | 8-bit · 17 tok/s |

## HOW TO RUN IT

| | |
|---|---|
| LLAMA.CPP | `llama-server -hf unsloth/Qwen3-VL-8B-Thinking-GGUF` |
| VLLM | `vllm serve Qwen/Qwen3-VL-8B-Thinking` |
| MODEL CARD | huggingface.co/Qwen/Qwen3-VL-8B-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-1SK8Q5H, generated November 22, 2026.