GPUpartner

VALIDATION REPORT · TEXT GENERATION

# Mistral Small 3.2 24B

on RTX 3060 12 GB · 4-bit (Q4_K_M) · Ollama / LM Studio · Chat

STATUS · OCT 22, 2026

**Not validated**

0 of 2 checks pass

**Mistral Small 3.2 24B at 4-bit needs 15.7 GB; the RTX 3060 12 GB can use 12.0 GB. The smallest machine we checked that runs it is the RTX 4090.**

## VALIDATION

| CHECK | TARGET | RESULT | STATUS |
|---|---|---|---|
| Fits in GPU memory | 12.0 GB usable | 15.7 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 |

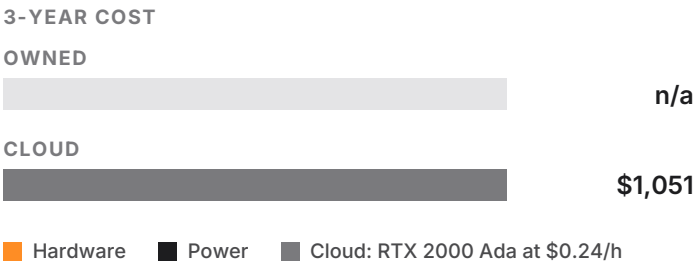


Budget basis: 4 h a day, power at $0.18 per kWh, cloud at RunPod on-demand prices as of Oct 22, 2026. Hardware at the MSRP entered in the calculator, not a GPU Partner quote. A 3-year view; resale value, staff time and taxes aren't counted.

GPUpartner

TECHNICAL DETAIL · TEXT GENERATION

# Mistral Small 3.2 24B

on RTX 3060 12 GB · 4-bit (Q4_K_M) · Ollama / LM Studio · Chat

## SETUP

| | |
|---|---|
| PARAMETERS | 24B |
| PRECISION | 4-bit (Q4_K_M) |
| CONTEXT | 8K, auto |
| CHATS AT ONCE | 1 |
| USE | Chat |
| RUNS IN | Ollama / LM Studio |
| HARDWARE | RTX 3060 12 GB |
| GPU CAN USE | 12.0 GB of 12.0 GB |

## MEMORY

| | |
|---|---|
| WEIGHTS | 13.6 GB |
| CONTEXT | 1.3 GB |
| RUNTIME | 877 MB |
| SYSTEM | 0.0 GB |
| OVER BY | 3.7 GB |

## SUGGESTED CHANGES

| | |
|---|---|
| 2× RTX 3060 12 GB | Fits across the cards, 48K context |
| SPILL TO RAM | Runs, slower, 7.8 tok/s |
| RTX 4090 | Runs it with 52K context |
| M5 MAX · 36 GB | Runs it with 52K context |

## SAME MODEL, OTHER MACHINES

| | |
|---|---|
| RTX 4090 | 6-bit · 35 tok/s |
| RTX 5090 | 8-bit · 48 tok/s |
| RTX PRO 6000 BLACKWELL | 8-bit · 48 tok/s |
| M5 MAX 128 GB | 8-bit · 18 tok/s |
| M3 ULTRA 256 GB | 8-bit · 20 tok/s |
| DGX SPARK 128 GB | 8-bit · 7.6 tok/s |
| RYZEN AI MAX+ 395 128 GB | 8-bit · 6.1 tok/s |

## HOW TO RUN IT

| | |
|---|---|
| OLLAMA | `ollama run mistral-small3.2:24b` |
| LLAMA.CPP | `llama-server -hf unsloth/Mistral-Small-3.2-24B-Instruct-2506-GGUF` |
| VLLM | `vllm serve mistralai/Mistral-Small-3.2-24B-Instruct-2506` |
| MLX (MAC) | `mlx_lm.chat --model mlx-community/Mistral-Small-3.2-24B-Instruct-2506-4bit` |
| MODEL CARD | huggingface.co/mistralai/Mistral-Small-3.2-24B-Instruct-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 Oct 2026; cloud prices are RunPod's on-demand list as of Oct 22, 2026. Report VR-11NJP8J, generated October 22, 2026.