GPUpartner

VALIDATION REPORT · TEXT GENERATION

Mistral Small 4 119B

on Radeon AI PRO R9700 · 4-bit (Q4_K_M) · Ollama / LM Studio · Chat

STATUS · NOV 22, 2026

Not validated

0 of 2 checks pass

Mistral Small 4 119B at 4-bit needs 68.4 GB; the Radeon AI PRO R9700 can use 32.0 GB. The smallest machine we checked that runs it is the RTX PRO 6000 Blackwell.

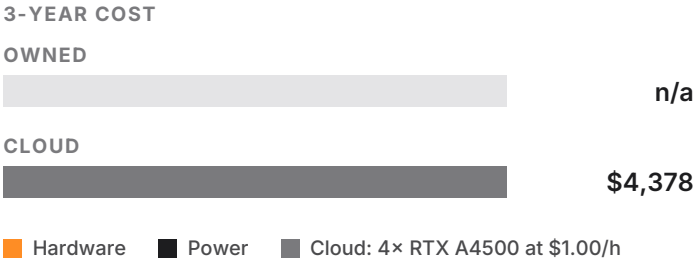
VALIDATION

CHECK	TARGET	RESULT	STATUS
Fits in GPU memory	32.0 GB usable	68.4 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	$1,459
3-YEAR TOTAL, OWNED	n/a
3-YEAR TOTAL, CLOUD	$4,378
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

Mistral Small 4 119B

on Radeon AI PRO R9700 · 4-bit (Q4_K_M) · Ollama / LM Studio · Chat

SETUP

PARAMETERS	119B, 6.5B active
PRECISION	4-bit (Q4_K_M)
CONTEXT	8K, auto
CHATS AT ONCE	1
USE	Chat
RUNS IN	Ollama / LM Studio
HARDWARE	Radeon AI PRO R9700
GPU CAN USE	32.0 GB of 32.0 GB

MEMORY

WEIGHTS	67.4 GB
CONTEXT	180 MB
RUNTIME	845 MB
SYSTEM	0.0 GB
OVER BY	36.4 GB

SUGGESTED CHANGES

3× RADEON AI PRO R9700	Fits across the cards, 256K context
SPILL TO RAM	Runs, slower, 23 tok/s
RTX PRO 6000 BLACKWELL	Runs it with 256K context
M5 MAX · 128 GB	Runs it with 256K context

SAME MODEL, OTHER MACHINES

RTX 4090	Too big
RTX 5090	Too big
RTX PRO 6000 BLACKWELL	5-bit · 171 tok/s
M5 MAX 128 GB	5-bit · 66 tok/s
M3 ULTRA 256 GB	FP8 · 57 tok/s
DGX SPARK 128 GB	FP8 · 27 tok/s
RYZEN AI MAX+ 395 128 GB	5-bit · 30 tok/s

HOW TO RUN IT

LLAMA.CPP	`llama-server -hf unsloth/Mistral-Small-4-119B-2603-GGUF`
VLLM	`vllm serve mistralai/Mistral-Small-4-119B-2603`
MLX (MAC)	`mlx_lm.chat --model mlx-community/Mistral-Small-4-119B-2603-4bit`
MODEL CARD	huggingface.co/mistralai/Mistral-Small-4-119B-2603

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-1UZ4LKD, generated November 22, 2026.