

VALIDATION REPORT · TEXT GENERATION

Nemotron 3 Super 120B A12B

on RTX 4060 Ti 16 GB · 4-bit (Q4_K_M) · Ollama / LM Studio · Chat

STATUS · OCT 22, 2026

Not validated

0 of 2 checks pass

Nemotron 3 Super 120B A12B at 4-bit needs 70.8 GB; the RTX 4060 Ti 16 GB can use 16.0 GB. The smallest machine we checked that runs it is the RTX PRO 6000 Blackwell.

VALIDATION

CHECK	TARGET	RESULT	STATUS
Fits in GPU memory	16.0 GB usable	70.8 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

3-YEAR COST

OWNED



n/a

CLOUD



\$4,378

Hardware Power Cloud: 4× RTX A4500 at \$1.00/h

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.

TECHNICAL DETAIL · TEXT GENERATION

Nemotron 3 Super 120B A12B

on RTX 4060 Ti 16 GB · 4-bit (Q4_K_M) · Ollama / LM Studio · Chat

SETUP

PARAMETERS	124B, 12B active
PRECISION	4-bit (Q4_K_M)
CONTEXT	8K, auto
CHATS AT ONCE	1
USE	Chat
RUNS IN	Ollama / LM Studio
HARDWARE	RTX 4060 Ti 16 GB
GPU CAN USE	16.0 GB of 16.0 GB

MEMORY

WEIGHTS	69.8 GB
CONTEXT	226 MB
RUNTIME	845 MB
SYSTEM	0.0 GB
OVER BY	54.8 GB

SUGGESTED CHANGES

SPILL TO RAM	Runs, slower, 9.5 tok/s
RTX PRO 6000 BLACKWELL	Runs it with 256K context
M5 MAX · 128 GB	Runs it with 256K context
RYZEN AI MAX+ 395 · 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 · 82 tok/s
M5 MAX 128 GB	5-bit · 32 tok/s
M3 ULTRA 256 GB	8-bit · 27 tok/s
DGX SPARK 128 GB	6-bit · 17 tok/s
RYZEN AI MAX+ 395 128 GB	5-bit · 15 tok/s

HOW TO RUN IT

OLLAMA	ollama run nemotron-3-super:120b-a12b
LLAMA.CPP	llama-server -hf unsloth/NVIDIA-Nemotron-3-Super-120B-A12B-GGUF
VLLM	vllm serve nvidia/NVIDIA-Nemotron-3-Super-120B-A12B-BF16
MLX (MAC)	mlx_lm.chat --model mlx-community/NVIDIA-Nemotron-3-Super-120B-A12B-4bit
MODEL CARD	huggingface.co/nvidia/NVIDIA-Nemotron-3-Super-120B-A12B-BF16

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-1YWFG6Q, generated October 22, 2026.