GPUpartner

VALIDATION REPORT · VIDEO GENERATION

# LTX-2.5

on H100 SXM · 16-bit (BF16)

STATUS · OCT 22, 2026

**Validated**

3 of 3 checks pass

**LTX-2.5 at 16-bit runs on the H100 SXM, about 26 s per clip.**

## VALIDATION

| CHECK | TARGET | RESULT | STATUS |
|---|---|---|---|
| Fits in GPU memory | 80.0 GB usable | 40.4 GB needed | PASS |
| Headroom | 5% free | 39.6 GB free (49%) | PASS |
| Time per clip | under 15 min | 26 s | PASS |

Targets: a clip in under 15 min, and 5% of memory left free. Calculated from published specs, not benchmarked.

## BUDGET

| | |
|---|---|
| HARDWARE, ONE-TIME | MSRP not entered |
| POWER, PER YEAR | $186 · 710 W at the wall |
| CLOUD INSTEAD, PER YEAR | $861 |
| 3-YEAR TOTAL, OWNED | Hardware + $559 power |
| 3-YEAR TOTAL, CLOUD | $2,583 |
| BREAK-EVEN | Needs the hardware MSRP |

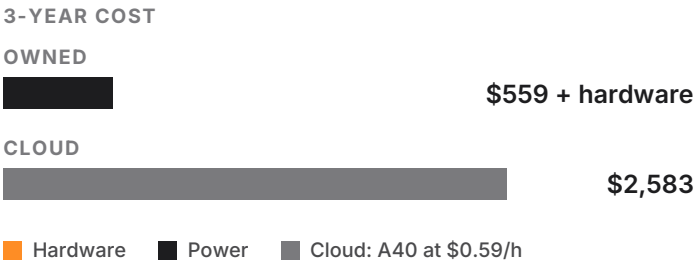


Budget basis: 4 h a day, power at $0.18 per kWh, cloud at RunPod on-demand prices as of Oct 22, 2026. Hardware at the MSRP entered in the calculator, not a GPU Partner quote. A 3-year view; resale value, staff time and taxes aren't counted.

GPUpartner

TECHNICAL DETAIL · VIDEO GENERATION

# LTX-2.5

on H100 SXM · 16-bit (BF16)

## SETUP

| | |
|---|---|
| PARAMETERS | 21B |
| PRECISION | 16-bit (BF16) |
| OUTPUT | 960 × 544, 121 frames |
| STEPS | 8 |
| HARDWARE | H100 SXM |
| GPU CAN USE | 80.0 GB of 80.0 GB |

## MEMORY

| | |
|---|---|
| WEIGHTS | 39.1 GB |
| WORKING | 542 MB |
| RUNTIME | 819 MB |
| SYSTEM | 0.0 GB |
| FREE | 39.6 GB |

## SAME MODEL, OTHER MACHINES

| | |
|---|---|
| RTX 4090 | 8-bit · 54 s |
| RTX 5090 | 8-bit · 47 s |
| RTX PRO 6000 BLACKWELL | 16-bit · 31 s |
| M5 MAX 128 GB | 16-bit · 2.1 min |
| M3 ULTRA 256 GB | 16-bit · 3.8 min |
| DGX SPARK 128 GB | 16-bit · 1.8 min |
| RYZEN AI MAX+ 395 128 GB | 16-bit · 3.1 min |

## WHERE TO GET IT

| | |
|---|---|
| FP8 | huggingface.co/unsloth/LTX-2.5-FP8 |
| GGUF | huggingface.co/unsloth/LTX-2.5-GGUF |
| OFFICIAL | huggingface.co/Lightricks/LTX-2.5 |

## METHOD AND ASSUMPTIONS

Memory is the denoiser's weights at the chosen precision plus the bigger of one transformer block's working set and the VAE decode, plus the runtime; the text encoders run first and then wait in system RAM. Time is the denoising FLOPs per step times the steps, against the chip's dense 16-bit TFLOPS, scaled to published GeForce runs of this model. 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-13SR8RK, generated October 22, 2026.