GPUpartner

VALIDATION REPORT · IMAGE GENERATION

# Stable Diffusion XL 1.0

on RTX 3080 10 GB · 16-bit (BF16)

STATUS · NOV 22, 2026

**Validated**

3 of 3 checks pass

**Stable Diffusion XL 1.0 at 16-bit runs on the RTX 3080 10 GB, about 11 s per image.**

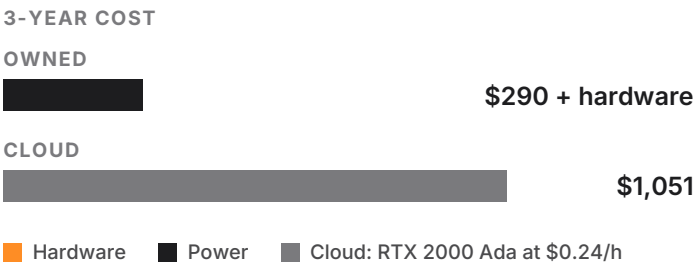
## VALIDATION

| CHECK | TARGET | RESULT | STATUS |
|---|---|---|---|
| Fits in GPU memory | 10.0 GB usable | 5.6 GB needed | PASS |
| Headroom | 5% free | 4.4 GB free (44%) | PASS |
| Time per image | under 30 s | 11 s | PASS |

Targets: an image in under 30 s, and 5% of memory left free. Calculated from published specs, not benchmarked.

## BUDGET

| | |
|---|---|
| HARDWARE, ONE-TIME | MSRP not entered |
| POWER, PER YEAR | $97 · 368 W at the wall |
| CLOUD INSTEAD, PER YEAR | $350 |
| 3-YEAR TOTAL, OWNED | Hardware + $290 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 · IMAGE GENERATION

# Stable Diffusion XL 1.0

on RTX 3080 10 GB · 16-bit (BF16)

## SETUP

| | |
|---|---|
| PARAMETERS | 2.6B |
| PRECISION | 16-bit (BF16) |
| OUTPUT | 1024 × 1024 |
| STEPS | 30 |
| HARDWARE | RTX 3080 10 GB |
| GPU CAN USE | 10.0 GB of 10.0 GB |

## MEMORY

| | |
|---|---|
| WEIGHTS | 4.8 GB |
| WORKING | 22 MB |
| RUNTIME | 819 MB |
| SYSTEM | 0.0 GB |
| FREE | 4.4 GB |

## SAME MODEL, OTHER MACHINES

| | |
|---|---|
| RTX 4090 | 16-bit · 4.7 s |
| RTX 5090 | 16-bit · 3.9 s |
| RTX PRO 6000 BLACKWELL | 16-bit · 2.2 s |
| M5 MAX 128 GB | 16-bit · 12 s |
| M3 ULTRA 256 GB | 16-bit · 23 s |
| DGX SPARK 128 GB | 16-bit · 11 s |
| RYZEN AI MAX+ 395 128 GB | 16-bit · 19 s |

## WHERE TO GET IT

| | |
|---|---|
| GGUF | huggingface.co/AEUPH/sdxl-unet-q8-gguf |
| OFFICIAL | huggingface.co/stabilityai/stable-diffusion-xl-base-1.0 |

## 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 Nov 2026; cloud prices are RunPod's on-demand list as of Nov 22, 2026. Report VR-1PQXTRY, generated November 22, 2026.