

SD.Next supported models

<https://vladmandic.github.io/sdnext-docs/Models/>

List of Models by Size

Created on 2025-07-15. Last update 2025-08-08

| Publisher | Model | Version | Size | Diffusion Architecture | Model Params | Text Encoder(s) | TE Params | Auto Encoder | |
|-------------------|-----------------------|-----------------------------|------------------|------------------------|--------------|--|-------------------------------|--------------|--|
| Wan-AI | WAN | 2.2 A14B | 118G | MMDiT | 14.28B | UMT5-XXL | 5.6B | 48ch VAE | |
| Wan-AI | WAN | 2.1 14B | 78.52GB | MMDiT | 14.28B | UMT5-XXL | 5.6B | 16ch VAE | |
| HiDream-AI | HiDream | I2 Fast/Dev/Full | 42.71 GB + 15.69 | MMDiT | 17.10B | CLiP ViT-L + ViT+G + T5-XXL + LLama-3.1-8B | 0.12B + 0.69B + 2.95B + 4.54B | 16ch VAE | |
| Qwen | Qwen-Image | 1 | 56.10GB | MMDiT | 20.43 | QWen-2.5 | 8.29B | | |
| nVidia | Cosmos - Predict2 T2I | 14B | 37.36GB | MMDiT | 14.26B | T5-XXL | 4.86 | WAN-VAE | |
| Black Forest Labs | Flux | 1 Dev/Sc hnell Kontext Krea | 32.93GB | MMDiT | 11.9B | CLiP ViT-L + T5-XXL | 0.12B + 4.76B | 16ch VAE | |
| Shuttle AI | shuttle | 3.1 | 32.93GB | MMDiT | 11.9B | CLiP ViT-L + T5-XXL | 0.12B + 4.76B | 16ch VAE | |
| WAN-AI | WAN | 2.2 5B | 32G | tbd | 5B | tbd | tbd | tbd | |

| Publisher | Model | Version | Size | Diffusion Architecture | Model Params | Text Encoder(s) | TE Params | Auto Encoder | |
|--------------------|-------------------------|-------------------|----------------|------------------------|--------------|------------------------------------|------------------------------|-----------------|--|
| FAL | AuraFlow | 0.3 | 31.90GB | MMDiT | 6.8B | UMT5 | 12.1B | VAE | |
| VectorSpaceLab | OmniGen | v2 | 30.50GB | Transformer | 3.97B | Qwen-VL-2.5 | 3.75B | VAE | |
| Thudm | CogView | 4 | 30.39GB | DiT | 6.37B | GLM-4 | 9.40B | VAE | |
| Kandinsky | Kandinsky | 3 | 27.72GB | Unet | 3.05B | T5-XXXL | 8.72B | VQ | |
| Wan-AI | WAN | 2.1 1.3B | 27.72GB | MMDiT | 1.42xB | UMT5-XXL | 5.68B | 16ch VAE | |
| StabilityAI | Stable Diffusion | 3.5 Large | 26.98GB | MMDiT | 8.05B | CLiP ViT-L + ViT+G + T5-XXL | 0.12B + 0.69B + 4.76B | 16ch VAE | |
| Iodestones | Chroma | 36 | 26.84GB | MMDiT | 8.9B | CLiP ViT-L + T5-XXL | 0.12B + 4.76B | 16ch VAE | |
| Ostris | Flex | 1 Alpha | 25.65GB | MMDiT | 4.0B | CLiP ViT-L + T5-XXL | 0.12B + 2.95B | 16ch VAE | |
| Thudm | CogView | 3 Plus | 24.96GB | DiT | 2.85B | T5-XXL | 4.76B | VAE | |
| PixArt | ?, ? | XL 2 | 21.3GB | DiT | 0.61B | T5-XXL | 4.76B | VAE | |
| AlphaV LLM | Lumina | 2 | 20.75GB | DiT | 2.61B | Gemma-2 | 2.61B | 16ch VAE | |
| FreePix | F-Lite | | 19.81GB | MMDiT | 9.8B | T5-XXL | 2.95B | 16ch VAE | |
| Bria AI | Bria | 3.2 | 18.66GB | MMDiT | 3.78B | T5-XXL | 4.76B | 16ch VAE | |
| Kwai | Kolors | N/A | 17.40GB | UNnet | 2.58B | ChatGLM | 6.24B | VAE | |
| StabilityAI | Stable Diffusion | 3.5 Medium | 15.89GB | MMDiT | 2.25B | CLiP ViT-L + ViT+G + T5-XXL | 0.12B + 0.69B + 4.76B | 16ch VAE | |

| Publisher | Model | Version | Size | Diffusion Architecture | Model Params | Text Encoder(s) | TE Params | Auto Encoder | |
|----------------|-----------------------|--------------|---------|------------------------|---------------|-----------------------------|-----------------------|--------------|--|
| NVLabs | Sana | 1.5 4.8B | 15.58GB | MMDiT | 4.72B | Gemma 2 | 2.61B | DC-AE | |
| DeepFloyd | IF | L | 15.48GB | Multi-stage UNet | 0.61B + 0.93B | T5-XXL | 4.76B | Pixel | |
| VectorSpaceLab | OmniGen | v1 | 15.47GB | Transformer | 3.76B | Phi-3 | 0 | VAE | |
| StabilityAI | Stable Diffusion | 3.0 Medium | 15.14GB | MMDiT | 2.0B | CLiP ViT-L + ViT+G + T5-XXL | 0.12B + 0.69B + 4.76B | 16ch VAE | |
| Tencent | HunyuanDiT | 1.2 | 14.09GB | DiT | 1.5B | BERT + T5-XL | 3.52B + 1.67B | VAE | |
| PlaygroundAI | Playground | 2.x | 13.35GB | UNet | 2.56B | CLiP ViT-L + ViT+G | 0.12B + 0.69B | VAE | |
| nVidia | Cosmos - Predict2 T2I | 2B | 13.32GB | MMDiT | 1.96B | T5-XXL | 4.86 | WAN-VAE | |
| DeepFloyd | IF | M | 12.79GB | Multi-stage UNet | 0.37B + 0.46B | T5-XXL | 4.76B | Pixel | |
| NVLabs | Sana | 1.0 1600M | 12.63GB | MMDiT | 1.60B | Gemma 2 | 2.61B | DC-AE | |
| Warp AI | Wuerstchen | N/A | 12.16GB | Multi-stage UNet | 1.0B + 1.05B | CLiP ViT-L + ViT+G | 0.12B + 0.69B | 42x VQE | |
| StabilityAI | Stable Cascade | Medium | 11.82GB | Multi-stage UNet | 1.56B + 3.6B | CLiP ViT-G | 0.69B | 42x VQE | |
| NVLabs | Sana | 1.5 1.6B | 9.49GB | MMDiT | 1.60B | Gemma 2 | 2.61B | DC-AE | |

| Publisher | Model | Version | Size | Diffusion Architecture | Model Params | Text Encoder(s) | TE Params | Auto Encoder | |
|--------------------|-------------------------|-----------|---------------|------------------------|--------------|---------------------------|----------------------|--------------|--|
| Segmind | SSD-1B | N/A | 8.72GB | UNet | 1.33B | CLiP ViT-L + ViT+G | 0.12B + 0.69B | VAE | |
| AlphaV LLM | Lumina | Next SFT | 8.67GB | DiT | 1.7B | Gemma | 2.5B | VAE | |
| NVLabs | Sana | 1.0 600M | 7.51GB | MMDiT | 0.59B | Gemma 2 | 2.61B | DC-AE | |
| Salesforce | BLIP-Diffusion | N/A | 7.23GB | UNet | 0.86B | CLiP ViT-L + BLiP-2 | 0.12B + 0.49B | VAE | |
| StabilityAI | Stable Diffusion | XL | 6.94GB | UNet | 2.56B | CLiP ViT-L + ViT+G | 0.12B + 0.69B | VAE | |
| Koala | Koala | 700M | 6.58GB | UNet | 0.78B | CLiP ViT-L + ViT+G | 0.12B + 0.69B | VAE | |
| Segmind | Vega | N/A | 6.43GB | UNet | 0.75B | CLiP ViT-L + ViT+G | 0.12B + 0.69B | VAE | |
| Thu-ML | UniDiffuser | v1 | 5.37GB | U-ViT | 0.95B | CLiP ViT-L + CLiP ViT-B | 0.12B + 0.16B | VAE | |
| Kandinsky | Kandinsky | 2.2 | 5.15GB | Unet | 1.25B | CLiP ViT-G | 0.69B | VQ | |
| Stability AI | Stable Cascade | Lite | 4.97GB | Multi-stage UNet | 0.7B + 1.0B | CLiP ViT-G | 0.69B | 42x VQE | |
| PlaygroundAI | Playground | 1 | 4.95GB | UNet | 0.86B | CLiP ViT-L | 0.12B | VAE | |
| MeissoonFlow | Meissoonic | N/A | 3.64GB | DiT | 1.18B | CLiP ViT-H | 0.35B | VQ | |
| Open-MUSE | aMUSED | 256 | 3.41GB | ViT | 0.60B | CLiP ViT-L | 0.12B | VQ | |
| Stability AI | Stable Diffusion | 2.1 | 2.58GB | UNet | 0.86B | CLiP ViT-H | 0.34B | VAE | |

| Publisher | Model | Version | Size | Diffusion Architecture | Model Params | Text Encoder(s) | TE Params | Auto Encoder | |
|-------------|------------------|---------|--------|------------------------|--------------|-----------------|-----------|--------------|--|
| StabilityAI | Stable Diffusion | 1.5 | 2.28GB | UNet | 0.86B | CLiP ViT-L | 0.12B | VAE | |
| IDKiro | SDXS | N/A | 2.05GB | UNet | 0.32B | CLiP ViT-L | 0.12B | VAE | |
| Segmind | Tiny | N/A | 1.03GB | UNet | 0.32B | CLiP ViT-L | 0.12B | VAE | |

Generation Examples

Searching for highest resolution available on SD.Next ipex (dev 2025-07-15) on ultra 9 185H + 128GB RAM

Test Prompt: VibrantlySharp style, 3pic_vist4, animeniji, in a soft-focus painterly concept art style Upper body profile view, autumn dryad elf adorned in translucent red foliage and bark-textured silk, wind animating her hair and dress into motion, background: ancient twilight forest, distant amber sunlight diffused through fog and drifting leaves.

Stable Diffusion 3.5 Large+ class models

Text to Image generation examples for models with size larger than 26.8GB. It is range from lodestone Chroma and SD3.5L which both are about 27GB in size. And up to Wan 2.1 14B which is about 75GB in size. Decent quality is generated in about 10 minutes on this Model size class on Intel Core Ultra 9 185H with 128GB of DDR5 (ipex, iGPU)

☐☐- Models require identification, access request, agree on terms

| Model | Low steps | High steps |
|-------|-----------|------------|
|-------|-----------|------------|

Wan2.2-T2V-A14B

256x256 - 2m 7.55s @ 24 STEPS

256x256 - 4m 21.77s @ 50 STEPS

512x512 - 5m 55.53s @ 20STEPS

768x768 - hangs

768x1024 - hangs

1024x1024 - 27m 6.16s @ 20STEPS

1024x1024 - 66m 59.70s @ 50STEPS

UniPC Multistep Scheduler

generation progress + preview

□□

2025-07-30

2025-07-31 - OOM for all



Parameters: Steps: 20 | Size: 1024x1024 | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: 636d261 | Pipeline: WanPipeline | Operations: txt2img | Model: Wan2.2-T2V-A14B-Diffusers

Time: 27m 6.16s | total 1721.93
pipeline 1625.83 preview 64.05 vae
10.38 offload 10.13 te 9.62 gc 2.15
post 0.31 | GPU 39894 MB 31% | RAM
29.77 GB 24%



Parameters: Steps: 50 | Size: 1024x1024 | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: 636d261 | Pipeline: WanPipeline | Operations: txt2img | Model: Wan2.2-T2V-A14B-Diffusers

Time: 66m 59.70s | total 4090.56
pipeline 4019.37 preview 46.47 vae
10.70 offload 10.06 gc 2.17 te 2.04
post 0.32 | GPU 40040 MB 31% | RAM
29.63 GB 24%

Wan2.1-T2V-14B

256x256 - 1m 41.11s @ 16 STEPS
512x512 - 5m 48.96s @ 20STEPS
512x512 - 11m 42.14s @ 40STEPS
512x768 - hangs

UniPC Multistep Scheduler

generation progress + preview



512x512- 5m 48.96s @ 20 STEPS



Parameters: Pipeline: WanPipeline|
Steps: 20| Size: 512x512| Seed:
1460879190| CFG scale: 6| App:
SD.Next| Version: 251c919|
Operations: txt2img| Model: Wan2.1-
T2V-14B-Diffusers

Time: 5m 48.96s | total 511.83
pipeline 348.65 preview 140.76
offload 9.90 vae 8.66 gc 2.12 te 1.99
post 0.29 | GPU 39054 MB 30% | RAM
29.72 GB 24%

512x512- 11m 42.14s @ 40 STEPS



Parameters: Pipeline: WanPipeline|
Steps: 40| Size: 512x512| Seed:
1460879190| CFG scale: 6| App:
SD.Next| Version: 251c919|
Operations: txt2img| Model: Wan2.1-
T2V-14B-Diffusers

Time: 11m 42.14s | total 1079.70
pipeline 701.83 preview 345.45 vae
12.42 offload 11.25 te 6.87 gc 2.12
post 0.28 | GPU 38830 MB 30% | RAM
29.44 GB 23%

HiDream I1 Full?

(Meta-Llama-3.1-8B-Instruct)

512x512-generates-1024x1024

768x768-generates-1024x1024

1024x1024 -10m 29.57s @ 12 STEPS

1024x1024 -16m 42.62s @ 20 STEPS

1024x1024 -41m 42.55s @ 50 STEPS

1280x1280-generates-1024x1024

2048x2048-generates-1024x1024

UniPC Multistep Scheduler

generation progress + preview

1024x1024 - **10m 29.57s** @ **12 STEPS**



Parameters: Pipeline: HiDreamImagePipeline| **Steps: 12**| **Size: 1024x1024**| Seed: 1460879190| CFG scale: 6| App: SD.Next| Version: 3c66c35| Operations: txt2img| Model: HiDream-I1-Full

Time: **10m 29.57s** | total 685.95
pipeline 593.18 move 21.17 prompt
21.13 te 18.92 decode 14.95 offload
11.79 preview 4.42 gc 0.61 post 0.27
| GPU **4508** MB 4% | RAM **58.59** GB
47%

1024x1024 - **41m 42.55s** @ **50 STEPS**



Parameters: Pipeline: HiDreamImagePipeline| **Steps: 50**| **Size: 1024x1024**| Seed: 1460879190| CFG scale: 6| App: SD.Next| Version: 3c66c35| Operations: txt2img| Model: HiDream-I1-Full

Time: **41m 42.55s** | total 2576.47
pipeline 2466.40 preview 22.89 move
20.63 prompt 20.52 te 18.32 decode
15.24 offload 12.13 gc 0.55 post 0.27
| GPU **4572** MB 4% | RAM **58.59** GB
47%

HiDream I1 Dev?

1024x1024 **17m 47.83s** @ 20 STEP
1024x1024 **23m 59.26s** @ 28 STEP

FlowMatch LCM Scheduler

generation progress + preview

□□□□

1024x1024 - 17m 47.83s @ 20 STEPS



Parameters: Pipeline: HiDreamImagePipeline| **Steps: 20**| **Size: 1024x1024**| Seed: 1460879190| CFG scale: **2**| App: SD.Next| Version: 3c66c35| Operations: txt2img| Model: HiDream-I1-Dev

Time: **17m 47.83s** | total 1151.56
pipeline 1019.03 move 30.23 prompt
30.11 te 27.23 decode 18.25 offload
15.44 preview 10.76 gc 0.74 post
0.32 | GPU **4378** MB 3% | RAM **58.68**
GB 47%

1024x1024 - 23m 59.26s @ 28 STEPS



Parameters: Pipeline: HiDreamImagePipeline| **Steps: 28**| **Size: 1024x1024**| Seed: 1460879190| CFG scale: **2**| App: SD.Next| Version: 3c66c35| Operations: txt2img| Model: HiDream-I1-Dev

Time: **23m 59.26s** | total 1529.78
pipeline 1389.13 move 34.36 prompt
34.23 offload 22.57 te 21.92 decode
15.51 preview 11.74 gc 0.53 post
0.26 | GPU **4510** MB 4% | RAM **58.7**
GB 47%

HiDream I1 Fast?

CFG (0-2)

1024x1024 3m 39.03s @ 4 STEP


1024x1024 20m 16.60s @ 24 STEP

FlowMatch LCM Scheduler

generation progress + preview

|||||

1024x1024 - 3m 39.03s @ 4 STEPS



Parameters: Pipeline: HiDreamImagePipeline| Steps: 4| Size: 1024x1024| Seed: 1460879190| CFG scale: 2| App: SD.Next| Version: 3c66c35| Operations: txt2img| Model: HiDream-I1-Fast

Time: 3m 39.03s | total 249.11
pipeline 194.59 offload 18.99 decode
14.74 move 9.39 prompt 9.24 preview
1.85 gc 0.57 post 0.32 | GPU 4440 MB
3% | RAM 58.61 GB 47%

1024x1024 - 20m 16.60s @ 24 STEPS



Parameters: Pipeline: HiDreamImagePipeline| Steps: 24| Size: 1024x1024| Seed: 1460879190| CFG scale: 2| App: SD.Next| Version: 3c66c35| Operations: txt2img| Model: HiDream-I1-Fast

Time: 20m 16.60s | total 1283.51
pipeline 1178.55 move 22.60 prompt
22.47 te 20.27 decode 15.17 offload
12.25 preview 11.87 gc 0.53 post
0.27 | GPU 4502 MB 4% | RAM 58.63
GB 47%

HiDream E1 Full ?

768x768 generates 1024x1024
1024x1024 7m 33.98s @ 8 STEP
1024x1024 10m 25.10s @ 12 STEP
1024x1024 13m 40.41s @ 16 STEP
1024x1024 26m 55.45s @ 32 STEP

UniPC Multistep Scheduler

generation progress + preview

□□□□

1024x1024 - **13m 40.41s** @ 16
STEPS



Parameters: Pipeline:
HiDreamImagePipeline| **Steps: 16**
Size: 1024x1024 | Seed: 1460879190|
CFG scale: **2** | App: SD.Next| Version:
3c66c35| Operations: txt2img| Model:
HiDream-E1-Full

Time: **13m 40.41s** | total 860.48
pipeline 782.43 move 21.89 prompt
19.66 decode 15.80 offload 12.73
preview 7.64 gc 0.52 post 0.27 | GPU
4500 MB 4% | RAM **58.7** GB 47%

1024x1024 - **26m 55.45s** @ 32
STEPS



Parameters: Pipeline:
HiDreamImagePipeline| **Steps: 32**
Size: 1024x1024 | Seed: 1460879190|
CFG scale: 6| App: SD.Next| Version:
3c66c35| Operations: txt2img| Model:
HiDream-E1-Full

Time: **26m 55.45s** | total 1665.14
pipeline 1578.31 move 20.86 prompt
18.67 preview 17.96 decode 16.01
offload 13.01 gc 0.54 post 0.27 | GPU
4564 MB 4% | RAM **58.64** GB 47%

Qwen-Image

With true_cfg=0

512x512 - 2m 0.22s @ 20 steps
768x768 - 5m 23.78s @ 20 steps
1024x1024 - 9m 18.46s @ 16 steps
1024x1024 - 11m 34.66s @ 20 steps
1024x1024 - 18m 46.91s @ 32 steps
1024x1536 - 21m 59.50s @ 20 steps
1600x1600 - 49m 24.79s @ 20 steps

With true_cfg=4

1024x1024 - 18m 46.05s @ 16 steps
1024x1024 - 37m 42.68s @ 32 steps


FlowMatch Euler DiscreteScheduler

progress without preview

□□□□

created 2025-08-05
update 2025-08-08


1024x1024 - 9m 18.46s @ 16 STEPS



Parameters: Steps: 16| Size: 1024x1024| Seed: 1460879190| CFG scale: 6| App: SD.Next| Version: 1d37a25| Pipeline: QwenImagePipeline| Operations: txt2img| Model: Qwen-Image

Time: 9m 18.46s | total 559.47
pipeline 558.44 te 0.98 | GPU 61110
MB 48% | RAM 3.1 GB 2%

1024x1024 - 18m 46.05s @ 16 STEPS



Parameters: Steps: 16| Size: 1024x1024| Seed: 1460879190| CFG scale: 6| CFG true: 4| App: SD.Next| Version: 9ef6baf| Pipeline: QwenImagePipeline| Operations: txt2img| Model: Qwen-Image

Time: 18m 46.05s | total 1132.26
pipeline 1125.98 te 4.41 vae 1.79 |
GPU 60262 MB 47% | RAM 3.13 GB 2%

1024x1024 - 18m 46.91s @ 32 STEPS



Parameters: Steps: 32| Size: 1024x1024| Seed: 1460879190| CFG scale: 6| App: SD.Next| Version: 1d37a25| Pipeline: QwenImagePipeline| Operations: txt2img| Model: Qwen-Image

Time: 18m 46.91s | total 1127.93
pipeline 1126.88 te 0.98 | GPU 61110
MB 48% | RAM 3.1 GB 2%

1024x1024 - 37m 42.68s @ 16 STEPS



Parameters: Steps: 32| Size: 1024x1024| Seed: 1460879190| CFG scale: 6| CFG true: 4| App: SD.Next| Version: 9ef6baf| Pipeline: QwenImagePipeline| Operations: txt2img| Model: Qwen-Image

Time: 37m 42.68s | total 2264.61
pipeline 2262.65 te 1.89 | GPU 60264
MB 47% | RAM 3.12 GB 2%

Cosmos-Predict2 T2I 14B ?

256x256 - 2m 31.30s @ 20 STEPS
512x512 - 5m 54.52s @ 20 STEPS
768x768 - 10m 23.35s @ 16 STEPS
768x768 - 12m 41.39s @ 20 STEPS
768x768 - 20m 9.11s @ 32 STEPS
768x1024 - hangs


FlowMatch Euler DiscreteScheduler

progress without preview

"diffusers_offload_mode": "model"

□□□□

768x768- 10m 23.35s @ 16 STEPS



Parameters: Steps: 16 | Size: 768x768 | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: af58803 | Pipeline: Cosmos2TextToImagePipeline | Operations: txt2img | Model: Cosmos-Predict2-14B-Text2Image

Time: 10m 23.35s | total 643.47
pipeline 623.02 te 11.87 vae 8.56
post 0.32 gc 0.31 | GPU 27938 MB
22% | RAM 38.86 GB 31%

768x768- 20m 9.11s @ 32 STEPS



Parameters: Steps: 32 | Size: 768x768 | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: af58803 | Pipeline: Cosmos2TextToImagePipeline | Operations: txt2img | Model: Cosmos-Predict2-14B-Text2Image

Time: 20m 9.11s | total 1224.37
pipeline 1208.85 vae 7.85 te 7.65
post 0.26 | GPU 27938 MB 22% | RAM
38.88 GB 31%

[shuttle-3.1-aesthetic](#)

512x512 - 6m 11.47s @ 50 STEPS
768x768 - 4m 35.14s @ 20 STEPS
768x1024 - 6m 3.73s @ 20 STEPS
1024x1024 - 8m 5.95s @ 20 STEPS
1024x1280 - 4m 31.69s @ 8 STEPS
1024x1536 - 3m 58.29s @ 6 STEPS
1024x1536 - 20m 29.93s @ 32 STEPS
1280x1600 - 13m 44.74s @ 8 STEPS
1536x2048 - 12m 45.33s @ 6 STEPS

FlowMatch Euler DiscreteScheduler

generation progress + preview

|||||

2025-07-30

1024x1536 - 3m 58.29s @ 6

STEPS



Parameters: Steps: 6 | Size:

1024x1536 | Seed: 1460879190 | CFG
scale: 6 | App: SD.Next | Version:
12d7246 | Pipeline: FluxPipeline |
Operations: txt2img | Model: shuttle-
3.1-aesthetic

Time: 3m 58.29s | total 244.61
pipeline 229.67 decode 8.30 preview
6.31 gc 0.55 post 0.30 | GPU 37402
MB 29% | RAM 3.21 GB 3%

1024x1536 - 20m 29.93s @ 32

STEPS



Parameters: Steps: 32 | Size:

1024x1536 | Seed: 1460879190 | CFG
scale: 6 | App: SD.Next | Version:
12d7246 | Pipeline: FluxPipeline |
Operations: txt2img | Model: shuttle-
3.1-aesthetic

Time: 20m 29.93s | total 1270.02
pipeline 1221.19 preview 36.27
decode 8.43 prompt 1.92 te 1.89 gc
0.55 post 0.31 | GPU 37066 MB 29% |
RAM 3.22 GB 3%

FLUX.1-dev ?

64x64 - 58.97s @ 16 STEPS

128x128 - 1m 0.43s @ 16 STEPS

256x256 - 1m 7.51s @ 16 STEPS

512x512 - 2m 0.41s @ 16 STEPS

768x768 - 3m 41.69s @ 16 STEPS

768x1024 - 4m 54.68s @ 16 STEPS

768x1024 - 14m 56.94s @ 40 STEPS

1024x1536 - 5m 9.01s @ 8 STEPS

1024x1536 - 20m 13.60s @ 32 steps

FlowMatch Euler DiscreteScheduler

generation progress + preview

□□□□

Updated 2025-07-30

1024x1536 - 5m 9.01s @ 8 STEPS




Parameters: Steps: 8 | Size:

1024x1536 | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: 12d7246 | Pipeline: FluxPipeline | Operations: txt2img | Model: FLUX.1-dev

Time: 5m 9.01s | total 321.00 pipeline 300.51 decode 8.20 preview 8.10 prompt 1.96 te 1.94 gc 0.52 post 0.30 | GPU 37048 MB 29% | RAM 3.24 GB 3%

1024x1536 - 20m 13.60s @ 32 STEPS



Parameters: Steps: 32 | Size:

1024x1536 | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: 12d7246 | Pipeline: FluxPipeline | Operations: txt2img | Model: FLUX.1-dev

Time: 20m 13.60s | total 1217.57 pipeline 1204.89 decode 8.40 prompt 1.91 te 1.89 gc 0.54 post 0.31 | GPU 37048 MB 29% | RAM 3.23 GB 3%


FLUX.1-Krea-dev ?

512x512 - 7m 27.95s @ 20 STEPS
768x768 - 2m 24.02s @ 8 STEPS
768x1024 - 2m 23.60s @ 8 STEPS
1024x1024 - 6m 22.68s @ 16 STEPS
1024x1024 - 12m 35.18s @ 32steps
1024x1536 - hangs

FlowMatch Euler DiscreteScheduler

□□□□

Added 2025-08-01



Parameters: Steps: 16 | Size: 1024x1024 | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: 42706fb | Pipeline: FluxPipeline | Operations: txt2img | Model: FLUX.1-Krea-dev

Time: 6m 22.68s | total 471.62
pipeline 376.23 preview 84.90 decode
6.43 prompt 1.90 te 1.88 gc 0.25 |
GPU 34688 MB 27% | RAM 3.39 GB
3%



Parameters: Steps: 32 | Size: 1024x1024 | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: 42706fb | Pipeline: FluxPipeline | Operations: txt2img | Model: FLUX.1-Krea-dev

Time: 12m 35.18s | total 1494.92
pipeline 748.48 preview 731.88
decode 6.66 prompt 3.78 te 3.76 gc
0.31 | GPU 35284 MB 27% | RAM 2.97
GB 2%

FLUX.1-schnell ?

768x768 - ok

768x1024 - 1m 24.86s @ 4 STEPS

768x1024 - 7m 45.92s @ 24 STEPS

1024x1024 - hangs

FlowMatch Euler DiscreteScheduler

□□□□


768x1024 - 1m 24.86s @ 4 STEPS



Parameters: Pipeline: FluxPipeline|
Steps: 4| Size: 768x1024| Seed:
1460879190| CFG scale: 6| App:
SD.Next| Version: 251c919|
Operations: txt2img| Model: FLUX.1-
schnell

Time: 1m 24.86s | total 127.77
pipeline 74.92 preview 36.64 decode
9.63 offload 6.53 post 0.29 gc 0.27 |
GPU 33784 MB 26% | RAM 25.14 GB
20%

768x1024 - 7m 45.92s @ 24 STEPS



Parameters: Pipeline: FluxPipeline|
Steps: 24| Size: 768x1024| Seed:
1460879190| CFG scale: 6| App:
SD.Next| Version: 251c919|
Operations: txt2img| Model: FLUX.1-
schnell

Time: 7m 45.92s | total 903.61
pipeline 438.91 preview 400.19 move
14.96 prompt 14.96 te 14.91 decode
11.74 offload 7.80 post 0.30 gc 0.28 |
GPU 33614 MB 26% | RAM 25.17 GB
20%

[AuraFlow v0.3](#)

512x512 — 2m 4.61s @ STEPS 16

768x768 - 4m 37.98s @ STEPS 16

768x768 - 14m 17.12s @ STEPS 50

768x768 - 29m 46.26s @ STEPS 100

1024x1024 - hangs

FlowMatch Euler DiscreteScheduler

generation progress + preview

□

768x768- 14m 17.12s @ 50 STEPS



Parameters: Steps: 50 | Size:

768x768 | Seed: 1460879190 | CFG scale: 3.5 | App: SD.Next | Version: b61118f | Pipeline: AuraFlowPipeline | Operations: txt2img | Model: AuraFlow-v0.3

Time: 14m 17.12s | total 1163.36
pipeline 853.70 preview 306.23
decode 3.15 gc 0.47 post 0.25 | GPU
17638 MB 14% | RAM 2.75 GB 2%

768x768- 29m 46.26s @ 100 STEPS



Parameters: Steps: 100 | Size:

768x768 | Seed: 1460879190 | CFG scale: 3.5 | App: SD.Next | Version: b61118f | Pipeline: AuraFlowPipeline | Operations: txt2img | Model: AuraFlow-v0.3

Time: 29m 46.26s | total 3499.63
pipeline 1782.79 preview 1713.40
decode 3.13 gc 0.56 post 0.32 | GPU
17638 MB 14% | RAM 2.76 GB 2%

Wan2.2-TI2V-5B

~~256x256 - 1m 26.51s @ STEPS 50~~

~~512x512 - 47.81s @ STEPS 20~~

~~512x512 - 1m 48.58s @ STEPS 50~~

768x768 - 3m 50.42s @ STEPS 50

704x1280 - **2m 17.40s** @ STEPS 20

704x1280 - **5m 29.53s** @ STEPS 50

1024x1024 - 2m 28.14s @ STEPS 20


UniPC Multistep Scheduler

progress without preview



Parameters: Steps: 20 | Size: 704x1280 | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: 12d7246 | Pipeline: WanPipeline | Operations: txt2img | Model: Wan2.2-TI2V-5B-Diffusers

Time: **2m 17.40s** | total 140.87
pipeline 137.10 te 2.03 vae 1.68 post
0.29 gc 0.27 | GPU 25558 MB 20% |
RAM 3.09 GB 2%



Parameters: Steps: 50 | Size: 704x1280 | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: 12d7246 | Pipeline: WanPipeline | Operations: txt2img | Model: Wan2.2-TI2V-5B-Diffusers

Time: **5m 29.53s** | total 332.80
pipeline 329.25 te 2.02 vae 1.47 post
0.28 gc 0.25 | GPU 25558 MB 20% |
RAM 3.09 GB 2%

OmniGen2


512x512 - 2m 12.11s @ 20STEPS
768x768 - 5m 39.35s @ 20STEPS
1024x1024 - 9m 33.83s @ 16 STEPS
1024x1024 - 12m 35.39s @ 20STEPS
1024x1024 - 30m 29.41s @ 50STEPS
1280x1280 - 9m 30.35s @ 8STEPS
1280x1280 - 14m 26.83s @ 12STEPS
1280x1280 - 23m 49.50s @ 20STEPS
1280x1280 - 60m 14.48s @ 50STEPS
1600x1600 - 50m 50.81s @ 20STEPS
2048x2048 - hangs

FlowMatch Euler DiscreteScheduler

generation progress not shown

□□□□

1024x1024 - 9m 33.83s @ 16 STEPS



Parameters: Pipeline:

OmniGen2Pipeline| Steps: 16| Size: 1024x1024| Seed: 1460879190| CFG scale: 6| App: SD.Next| Version: 54bdf01| Operations: txt2img| Model: OmniGen2

Time: 9m 33.83s | total 573.61
pipeline 573.55 post 0.26 | GPU 18
054 MB 14% | RAM 2.71 GB 2%

1024x1024 - 30m 29.41s @ 50 STEPS



Parameters: Pipeline:

OmniGen2Pipeline| Steps: 50| Size: 1024x1024| Seed: 1460879190| CFG scale: 6| App: SD.Next| Version: 54bdf01| Operations: txt2img| Model: OmniGen2

Time: 30m 29.41s | total 1829.18
pipeline 1829.12 post 0.27 | GPU 18
054 MB 14% | RAM 2.72 GB 2%

CogView 4

256x256 - 1m 11.07s @ 16 STEPS
512x512 - 2m 44.92s @ 16 STEPS
768x768 - 6m 44.86s @ 16 STEPS
1024x1024 - 14m 9.97s @ 16 STEPS
1024x1024 - 43m 56.72s @ 50 STEPS
1280x1280 - hangs

FlowMatch Euler DiscreteScheduler

progress without preview

□□□

1024x1024 - 14m 9.97s @ 16 STEPS



Parameters: Steps: 16 | Size: 1024x1024 | Seed: 1460879190 | CFG scale: 3.5 | App: SD.Next | Version: fe61b5e | Pipeline: CogView4Pipeline | Operations: txt2img | Model: CogView4-6B

Time: 14m 9.97s | total 857.86
pipeline 825.42 decode 16.59 move
15.57 gc 0.53 post 0.29 | GPU 7926
MB 6% | RAM 31.98 GB 26%

1024x1024 - 43m 56.72s @ 50 STEPS



Parameters: Steps: 50 | Size: 1024x1024 | Seed: 1460879190 | CFG scale: 3.5 | App: SD.Next | Version: fe61b5e | Pipeline: CogView4Pipeline | Operations: txt2img | Model: CogView4-6B

Time: 43m 56.72s | total 2644.35
pipeline 2612.41 decode 16.72 move
14.90 gc 0.53 post 0.26 | GPU 7926
MB 6% | RAM 31.99 GB 26%

Kandinsky 3

768x768 - 1m 5.40s @ 20STEPS
768x1024 - 45.45s @ 8 STEP
768x1024 - 1m 9.35s @ 20 STEP
768x1024 - 2m 35.93s @ 64 STEP
1024x1024 - hangs

DDPM Scheduler

progress without preview



768x1024 - 1m 9.35s @ 20 STEPS




Parameters: Steps: 20| Size:

768x1024| Seed: 1460879190| CFG
scale: 6| App: SD.Next| Version:
fe61b5e| Pipeline:
Kandinsky3Pipeline| Operations:
txt2img| Model: kandinsky-3

Time: 1m 9.35s | total 73.34 pipeline
64.79 offload 4.26 decode 4.24 post
0.29 gc 0.27 | GPU 29500 MB 23% |
RAM 18.93 GB 15%

768x1024 - 2m 35.93s @ 64 STEPS



Parameters: Steps: 64| Size:

768x1024| Seed: 1460879190| CFG
scale: 6| App: SD.Next| Version:
fe61b5e| Pipeline:
Kandinsky3Pipeline| Operations:
txt2img| Model: kandinsky-3

Time: 2m 35.93s | total 159.72
pipeline 151.59 offload 4.05 decode
4.03 post 0.29 gc 0.27 | GPU 29482
MB 23% | RAM 18.93 GB 15%

Wan2.1-T2V-1.3B

256x256 - 23.16s @ 16 STEPS

512x512 - 54.41s @ 20STEPS

768x768 - 3m 43.11s @ 40STEPS

1024x1024 - 7m 35.17s @ 40STEPS

1280x1280 - 12m 42.76s @ 40STEPS

(resolution to be divisible by 16)

UniPC Multistep Scheduler

generation progress + preview

□

512x512 - 54.41s @ 20 STEPS



Parameters: Pipeline: WanPipeline|
Steps: 20| Size: 512x512| Seed:
1460879190| CFG scale: 6| App:
SD.Next| Version: 3c66c35|
Operations: txt2img| Model: Wan2.1-
T2V-1.3B-Diffusers

Time: 54.41s | total 96.96 pipeline
54.15 preview 36.26 gc 2.09 te 1.98
offload 1.88 vae 0.80 | GPU 15416 MB
12% | RAM 3.16 GB 3%

1280x1280- 12m 42.76s @ 40 STEPS



Parameters: Pipeline: WanPipeline|
Steps: 40| Size: 1280x1280| Seed:
1460879190| CFG scale: 6| App:
SD.Next| Version: 3c66c35|
Operations: txt2img| Model: Wan2.1-
T2V-1.3B-Diffusers

Time: 12m 42.76s | total 1282.37
pipeline 762.43 preview 511.77 vae
2.40 gc 2.11 te 1.98 offload 1.88 post
0.30 | GPU 21686 MB 17% | RAM 3.14
GB 3%

Stable Diffusion 3.5 L ?

512x512 - 4m 18.84s @ 32 STEPS
512x768 - 3m 25.31s @ 16 STEPS
512x768 - 11m 9.34s @ 50 STEPS
768x768 - @
1024x1024 - **14m 15.60s** @ 16 steps
1024x1024 - **28m 56.13s** @ 32 steps
1024x1536 - 29m 24.32s @ 16 steps
1024x1536 - 58m 47.43s @ 32 steps

FlowMatch Euler DiscreteScheduler

generation progress + preview

□□□□

update 2025-07-30



Parameters: Steps: 16 | Size: **1024x1024** | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: 12d7246 | Pipeline: StableDiffusion3Pipeline | Operations: txt2img | Model: stable-diffusion-3.5-large

Time: **14m 15.60s** | total 921.29 pipeline 848.32 preview 61.93 decode 6.98 prompt 3.64 gc 0.65 post 0.30 | GPU **29320** MB 23% | RAM 3.11 GB 2%



Parameters: Steps: 32 | Size: **1024x1024** | Sampler: DPM++ 2M SDE | Seed: 1460879190 | CFG scale: 6 | App: SD.Next | Version: b56d508 | Pipeline: StableDiffusion3Pipeline | Operations: txt2img | Model: stable-diffusion-3.5-large

Time: **28m 56.13s** | total 1743.07 pipeline 1729.54 preview 6.68 decode 6.57 gc 0.26 | GPU 29320 MB 23% | RAM 3.0 GB 2%

Chroma 36

512x512 - 4m 59.14s @ 20 STEPS

768x768 - 7m 4.50s @ 16 STEPS

768x768 - 21m 56.62s @ 50 STEPS

768X1024 - 12m 0.43s @ 20 steps

1024x1024 - hangs

progress with preview

FlowMatch Euler DiscreteScheduler

☐☐

only some CFG/resolutions is ok

768x768 - 7m 4.50s @ 16 STEPS



Parameters: Steps: 16 | Size:

768x768 | Seed: 1460879190 | CFG scale: 2 | App: SD.Next | Version: fe61b5e | Pipeline: ChromaPipeline | Operations: txt2img | Model: Chroma

Time: 7m 4.50s | total 627.03 pipeline 416.87 preview 197.83 decode 7.35 offload 4.69 gc 0.49 post 0.26 | GPU 11132 MB 9% | RAM 19.33 GB 15%

768x768 - 21m 56.62s @ 50 STEPS



Parameters: Steps: 50 | Size:

768x768 | Seed: 1460879190 | CFG scale: 2 | App: SD.Next | Version: fe61b5e | Pipeline: ChromaPipeline | Operations: txt2img | Model: Chroma

Time: 21m 56.62s | total 1945.92 pipeline 1308.96 preview 624.37 decode 7.38 offload 4.90 gc 0.50 post 0.27 | GPU 11154 MB 9% | RAM 19.62 GB 16%


chroma-unlocked-v48-detail-calibrated

512x768 - 4m 51.82s @ 16 STEPS
768x768 - 7m 6.56s @ 16 STEPS
768x768 - 14m 11.86s @ 32 STEPS
768x1024 - 9m 30.73s @ 16 STEPS
1024x1024 - 13m 3.35s @ 16 STEPS

FlowMatch Euler DiscreteScheduler




Added 2025-08-01



Parameters: Steps: 16 | Size: 768x768 | Seed: 1460879190 | CFG scale: 2 | App: SD.Next | Version: 42706fb | Pipeline: ChromaPipeline | Operations: txt2img | Model: chroma-unlocked-v48-detail-calibrated

Time: 7m 6.56s | total 563.91 pipeline 423.91 preview 137.06 decode 2.61 gc 0.28 | GPU 28122 MB 22% | RAM 2.94 GB 2%



Parameters: Steps: 32 | Size: 768x768 | Seed: 1460879190 | CFG scale: 2 | App: SD.Next | Version: 42706fb | Pipeline: ChromaPipeline | Operations: txt2img | Model: chroma-unlocked-v48-detail-calibrated

Time: 14m 11.86s | total 1162.50 pipeline 849.22 preview 302.91 prompt 3.72 te 3.70 decode 2.62 gc 0.30 | GPU 28134 MB 22% | RAM 2.95 GB 2%

SD.Next Config

```
{
  "samples_filename_pattern": "[seq]-[date]-[model_name]-[height]x[width]-Seed[seed]-CFG[cfg]-STEP[steps]",
  "diffusers_offload_mode": "none",
  "diffusers_to_gpu": true,
  "device_map": "gpu"
}
```

Dockerfile variables

```
# intel specific environment variables
ENV IPEX_SDPA_SLICE_TRIGGER_RATE=1
ENV IPEX_ATTENTION_SLICE_RATE=0.5
ENV IPEX_FORCE_ATTENTION_SLICE=1
ENV IPEXRUN=true
```

Revision #71

Created 2025-07-15 07:52:26 UTC by Oleksandr Liutyi

Updated 2025-08-08 14:36:44 UTC by Oleksandr Liutyi