

FREE GUIDE · FOR CREATORS

The Beginner's Guide to Local AI Image Generation

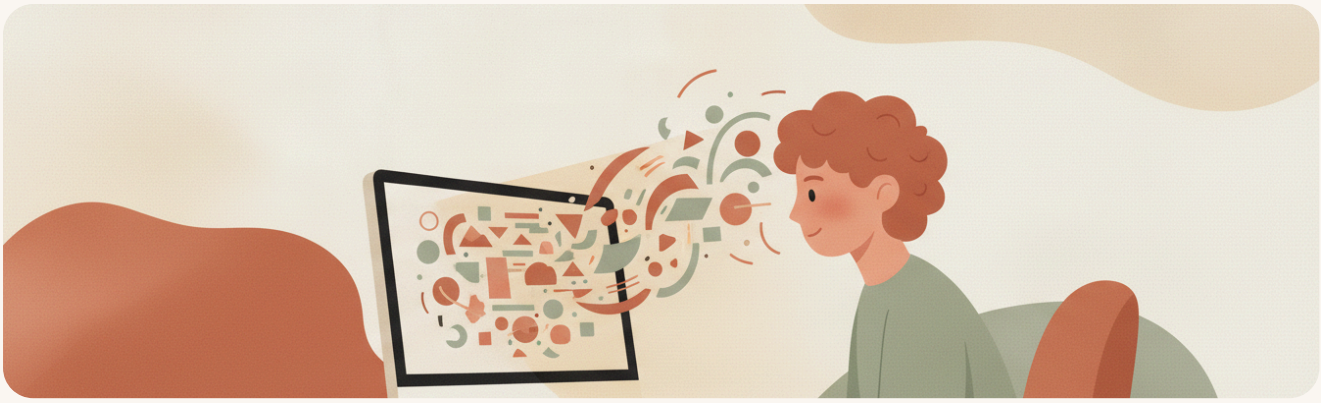
A practical introduction to generating images on your own machine — with realistic setup, cost and privacy considerations.

TITANXT

CONTENTS

What's Inside

01	**Why Local AI Is Taking Off**	The shift away from cloud services
02	**Cloud vs Local — What Actually Changes**	An honest side-by-side comparison
03	**What You Need to Get Started**	GPUs, RAM and the good news about gaming PCs
04	**Models Explained Without the Jargon**	Checkpoints, LoRAs and VAEs in plain English
05	**Prompting for Local AI Tools**	Structure, negative prompts, seeds
06	**Five Creative Workflow Ideas**	Products, social, thumbnails, concepts, ads
07	**Common Beginner Mistakes**	Six pitfalls and their fixes
08	**A Simpler Way to Work Locally**	Where Logen fits in
09	**Next Steps**	Start generating today



CHAPTER 01

Why Local AI Image Generation Is Taking Off

Something has shifted. While millions still type prompts into cloud services like Midjourney, DALL-E and Leonardo, a growing number of creators are going local — running AI image models directly on their own computers.

Why? Because local generation solves problems that cloud services create:

- **More creative control.**
  Local tools offer control over settings. Software restrictions, model licences and applicable rules still apply.

- **Different running costs.**
  Local generation can avoid cloud per-image fees. Hardware, electricity, storage and optional paid software still cost money.

- **More privacy control.**
  Local processing gives you more control over prompts and images. Downloads, software networking and optional services still matter.

- **Hardware-dependent speed.**
  Local processing avoids a shared cloud queue; loading and generation still take time and depend on hardware and settings.

- **Control and repeatability.**
  Save settings, seeds and model versions to repeat experiments. Hardware or software changes can alter results; identical output is not guaranteed.

> THE BIGGER PICTURE
>
> Some local tools require technical configuration. Accessible tools can simplify setup, but compatibility and results depend on the hardware, model and workflow.



CHAPTER 02

Cloud vs Local — What Actually Changes

Cloud generation

You type a prompt into a web interface, their hardware processes it, you receive an image.

- Subscription prices vary by provider
- Some services use per-image credits
- Provider-specific moderation policies
- Images stored on third-party servers
- Queues at peak hours; unannounced model swaps

Local generation

Models run on your own GPU. Prompts processed locally, images saved on your machine.

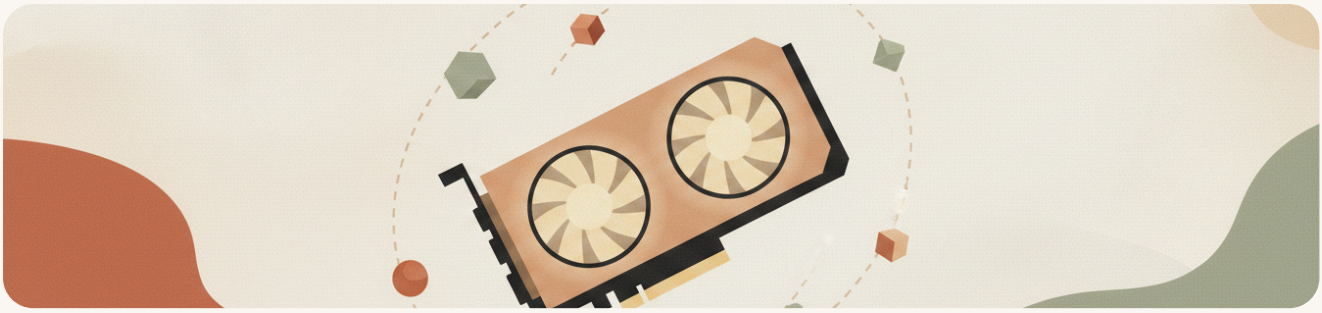
- Hardware, electricity and software costs remain
- Local processing; networking depends on software
- Save settings and seeds; outputs can vary
- Offline use depends on app requirements
- Needs a capable GPU + initial setup

The real-world comparison

Factor	Cloud (Midjourney / DALL-E)	Local AI
Monthly cost	Plan-dependent	Hardware, power, optional software
Per-image cost	Credit-based	No cloud fee; running costs remain
Privacy	Images on their servers	Generation local; networking varies
Content freedom	Filtered / moderated	Software rules and model terms apply
Speed	Queue-dependent	GPU-dependent
Consistency	Updates change outputs	You control the model
Offline use	Usually requires internet	App-dependent; see Pro checks in Ch. 08
Learning curve	Low	Medium (and improving)

IN PRACTICE

Compare local running costs and software requirements with your chosen cloud service. Either approach may suit different tasks.



CHAPTER 03

What You Need to Get Started

Hardware requirements vary by model and software. An important specification is **VRAM** — the GPU's dedicated memory.

Listed GPU requirement — 6 GB: not proof that every model runs at this capacity.

Documented generation check — 12 GB: NVIDIA RTX A2000; other GPUs are not verified here.

More VRAM: may help larger models and outputs; no universal speed guarantee.

Example GPUs — not verified Logen compatibility

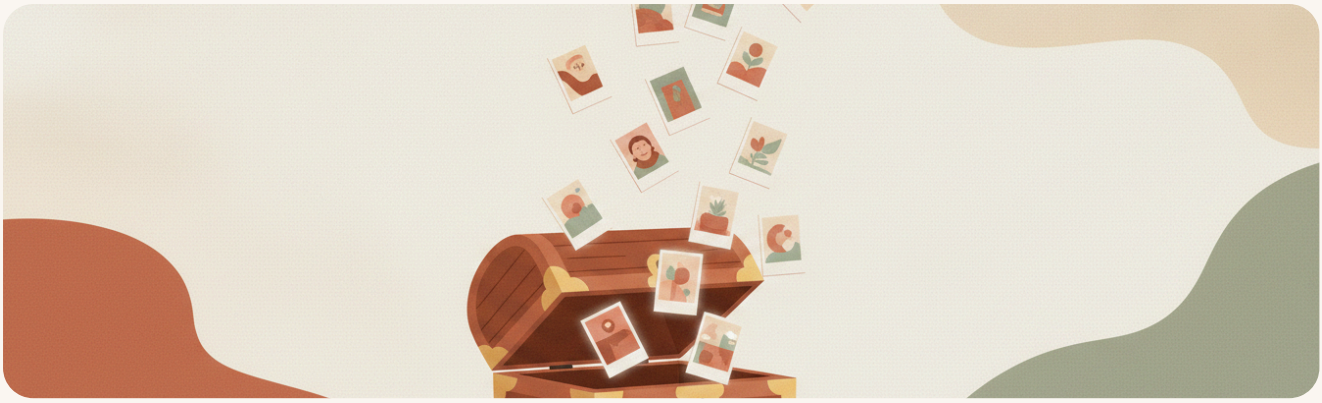
GPU	VRAM	Compatibility note
NVIDIA RTX 3060	12 GB	Check exact tool and model support
NVIDIA RTX 3070 / 3080	8–10 GB	Not verified for Logen here
NVIDIA RTX 4060 Ti	16 GB	Model and memory limits still apply
NVIDIA RTX 4070 / 4080	12–16 GB	Not verified for Logen here
NVIDIA RTX 4090	24 GB	Speed depends on model and settings
AMD RX 7900 XTX	24 GB	Do not assume Logen support

Check current release requirements — this general GPU list is not a Logen support guarantee. The documented test used an NVIDIA RTX A2000 with 12 GB VRAM.

RAM, storage and OS

- **RAM:** current Logen download guidance recommends 16 GB
- **Storage:** 30 GB free for the default Logen setup; additional models need more space
- **OS:** Logen: Windows and Linux x64. Other tools may support other platforms; this is not a macOS Logen recommendation.

> THE GOOD NEWS
>
> Check titanxt.io/download before installing. GPU memory, RAM, disk space, model and output size all affect what will run.



CHAPTER 04

Models Explained Without the Jargon

A model is a set of files (sizes vary widely) containing everything the AI has learned about generating images. Think of it like a camera lens: same body, different lens — different results.

The models you'll encounter

- **Stable Diffusion XL (SDXL)**
  — general local-AI context, not the Logen default. Typically used around 1024×1024; support varies by tool.

- **FLUX.1-schnell**
  — the Logen Free default. Similar model names do not imply the same memory needs, speed, workflows or licence.

- **SD 1.5**
  — an older model family used by other tools. Memory and feature support depend on the implementation.

Checkpoints, LoRAs and VAEs

Checkpoint — a complete model file; what people mean by "download a model." CivitAI and Hugging Face are the main sources. Examples for compatible tools (check exact licences): **Juggernaut XL** (photorealism), **DreamShaper XL** (versatile), **RealVisXL** (portraits).

LoRAs — small add-ons (10–200 MB) that teach a base model specific things: an art style, a character, a product, a lighting technique. Compatibility and combining support depend on the tool and base model.

VAEs — decode model representations into images. Packaging and compatibility vary; follow the exact model and workflow requirements.

> WHAT BEGINNERS ACTUALLY NEED
>
> In Logen, start with **FLUX.1-schnell**. Optional Pro models FLUX.2 [klein] 4B and Z-Image-Turbo have Apache-2.0 upstream bases; check exact weights and component terms. A purchase grants no blanket model or image rights. FLUX.1-dev and Qwen-Image-2.1 are disabled.



CHAPTER 05

Prompting for Local AI Tools

Prompt handling varies by model and tool. Treat the structure below as a starting point, not a universal rule or a quality guarantee.

Basic prompt structure

1. **Subject** — what's in the image (a woman, a landscape, a product)
2. **Style** — how it looks (photorealistic, oil painting, anime, minimalist)
3. **Details** — lighting, camera angle, time of day
4. **Optional descriptors** — detailed, sharp focus, professional; effects vary

> EXAMPLE
>
> "Product photograph of a ceramic coffee mug, white background, soft studio lighting, sharp focus, commercial photography, 4K, detailed texture"

Negative prompts

In models and workflows that support them, negative prompts describe what to avoid. They do not guarantee artifact-free results:

> EXAMPLE NEGATIVE
>
> "blurry, low quality, distorted, watermark, text, oversaturated, extra fingers, deformed"

Seeds and repeatability

Save the prompt, seed, model version and settings to make experiments easier to repeat. Software versions, hardware and settings can change the result; identical output is not guaranteed.

Practical tips

- **Be specific.** Describe the subject, setting and lighting you want; extra detail does not guarantee quality.
- **Prioritise** important elements clearly; position effects depend on the model and tool.
- **Separate concepts clearly** — commas can improve readability, not guarantee independent weighting.
- **Iterate.** Your first prompt won't be perfect. Adjust, regenerate, refine.

CHAPTER 06

Five Creative Workflow Ideas

1 Product mockups

1. General examples, not a Logen feature list; some steps need other tools. Photograph your product with even lighting
2. Remove and replace the background with AI
3. Generate studio, lifestyle or seasonal contexts
4. Export at your platform's required resolution

Why local matters: local processing offers control; review software networking before using sensitive designs.

2 Social media images

1. Define your brand style as a prompt template + LoRA combo
2. Generate images matching your content calendar
3. Keep consistent aspect ratios per platform
4. Generate variations as your tool supports; edit and assemble elsewhere if needed

Why local matters: saved settings support experiments; software and running costs still apply.

3 YouTube & blog thumbnails

1. Create a base template style (colours, composition, energy)
2. Generate the background or hero image
3. Leave space for text overlays
4. Maintain series consistency with saved settings

Why local matters: explore concepts at hardware-dependent speed; design and editing may need other tools.

4 Concept art & moodboards

1. Describe the concept in prompt form
2. Generate variations with different seeds
3. Combine favourites into a moodboard
4. Refine with collaborator feedback

Why local matters: iterate locally; review software and services before using sensitive ideas.

5 Ad creative testing

1. Define ad formats and safe zones
2. Generate multiple visual approaches per message
3. Export at platform specs and split-test
4. Double down on winners with more variants

Why local matters: explore variants without cloud per-image fees; speed and running costs vary.

CHAPTER 07

Common Beginner Mistakes

1 Installing too many tools

ComfyUI, Automatic1111, Fooocus and InvokeAI all at once — different formats, interfaces, learning curves.

Fix: Pick one tool. Learn it well. Expand later if needed.

2 Downloading every model

Dozens of CivitAI models before understanding any of them = confusion and wasted disk space.

Fix: In Logen, begin with FLUX.1-schnell. Learn one model before adding alternatives.

3 Ignoring image organisation

A week in, you'll have hundreds of images and no idea which prompt made which.

Fix: Use a tool that saves generation metadata. Project folders and meaningful names from day one.

4 Overcomplicating prompts

Long, unfocused prompts can be harder to refine; results depend on the model.

Fix: Start short and specific. Add detail only when the output needs it.

5 Wrong resolution for the model

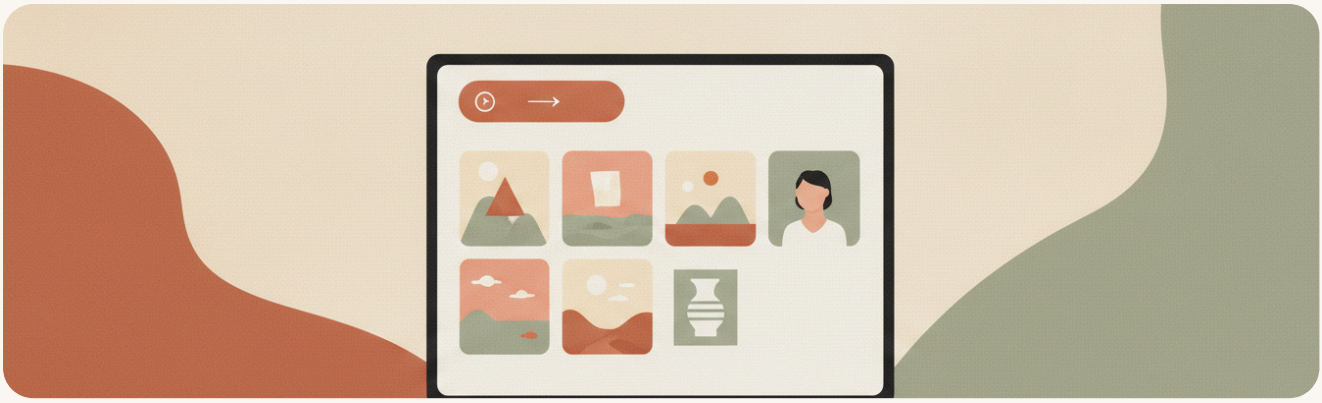
SDXL commonly starts at 1024×1024; SD 1.5 at 512×512. Check the exact model; other sizes may affect results.

Fix: Use your model's recommended resolution as baseline; upscale afterwards.

6 Giving up too early

Your first local images may differ from cloud results; some services modify prompts or apply extra processing.

Fix: Practise with focused comparisons. Improvement takes time; better-than-cloud results are not guaranteed.



CHAPTER 08

A Simpler Way to Work Locally

ComfyUI offers a node-based interface. Logen provides an interface for its supplied local-generation workflows without requiring you to build their node graphs; it does not promise every ComfyUI capability.

What if you didn't have to choose?

Logen is a desktop application built for creators who want local AI power without the complexity:

- **Platform-specific setup**
  — Windows/Linux x64; initial downloads and setup take time
- **Provided local workflows**
  — available models, features and results depend on the current release
- **Start with one model**
  — FLUX.1-schnell is Free; optional models have separate requirements and terms
- **Pro batch variations**
  — 2–16 seed variations of one prompt; not multi-prompt import or every editing pipeline
- **Local generation, online checks**
  — downloads need internet. Pro needs online activation and periodic entitlement checks; documented offline grace is 24 hours before fallback to Free. Local generation is not zero networking or absolute privacy.

Who Logen is for

- Designers who want AI as a creative tool, not a technical project
- Small business owners generating product and marketing imagery
- Content creators who need consistent, branded visuals
- Anyone frustrated by ComfyUI's complexity but unwilling to accept cloud limits

CHAPTER 09

Next Steps

You now understand the landscape — how local generation works, what you need, and the pitfalls to avoid. The next step is action.

- **Try Logen**
  — check titanxt.io/download for requirements. Setup includes downloads; generation time depends on hardware, model and settings.

- **Join the newsletter**
  — check available TitanXT emails and consent details before subscribing; no cadence or early-access promise is made here.

- **Explore the blog**
  — check available articles on workflows, models and creative techniques at titanxt.io/blog; no posting cadence is promised.

The creators adopting local AI today are building the skills and workflows that will define creative production for years. Cloud services will keep a place for quick exploration — compare control, networking, model terms and total costs for your own work.

Welcome to local AI. Your machine. Your experiments. Check the terms.

Plan your first local AI image

Logen is free to download for Windows and Linux x64. Check requirements, allow time for downloads, then start with FLUX.1-schnell. Pro requires online activation and periodic checks.

titanxt.io/download