Black Forest Labs

Founding Story and Product

Written out of my own interest and initiative · Gian Hossle · September 2026 · Sources are linked inline



Figure 1: Black Forest Labs at a glance. Sources: Crunchbase News, bfl.ai, Bloomberg, Sacra

Founding Story

Black Forest Labs was founded in 2024 by Robin Rombach (CEO), Andreas Blattmann (Cofounder) and Patrick Esser (Cofounder), together with a founding team of about ten researchers who had worked with them for years. That wider team includes Sumith Kulal and Dustin Podell, who both list themselves as founders, as well as Tim Dockhorn, Axel Sauer, Dominik Lorenz, Jonas Müller, Frederic Boesel and Harry Saini. The company is headquartered in Freiburg, Germany, and runs a second lab in San Francisco.

Few founding teams share a longer research history. Rombach, Esser and Blattmann did their early generative AI work in Björn Ommer's computer vision group, first at Heidelberg University and then at LMU Munich. In December 2020, Esser and Rombach published VQGAN, which taught a transformer to build images out of a compact "vocabulary" of image parts. A year later, Rombach, Blattmann, Lorenz and Esser introduced latent diffusion: instead of generating every pixel directly, the model works on a compressed version of the image, which significantly reduces the computing power needed for high resolution results. That idea became the core of Stable Diffusion, the open image model that set off the generative AI boom in 2022. Esser, who by then led applied research at Runway, described Stable Diffusion as "really the same model" that had come out of four years of work with Rombach and their collaborators in Ommer's group.

The group then moved to Stability AI, where it kept producing the company's most important models: SDXL in July 2023 (lead author Podell), Stable Video Diffusion (lead author Blattmann) and Adversarial Diffusion Distillation (lead author Sauer), a technique that produces high quality images in as few as four steps, in November 2023, and Stable Diffusion 3 in March 2024 (lead author Esser, with Kulal, Blattmann, Müller, Saini and Lorenz among the coauthors).

The founding team wrote the papers behind modern image generation

Authorship of six landmark papers by members of the Black Forest Labs founding team

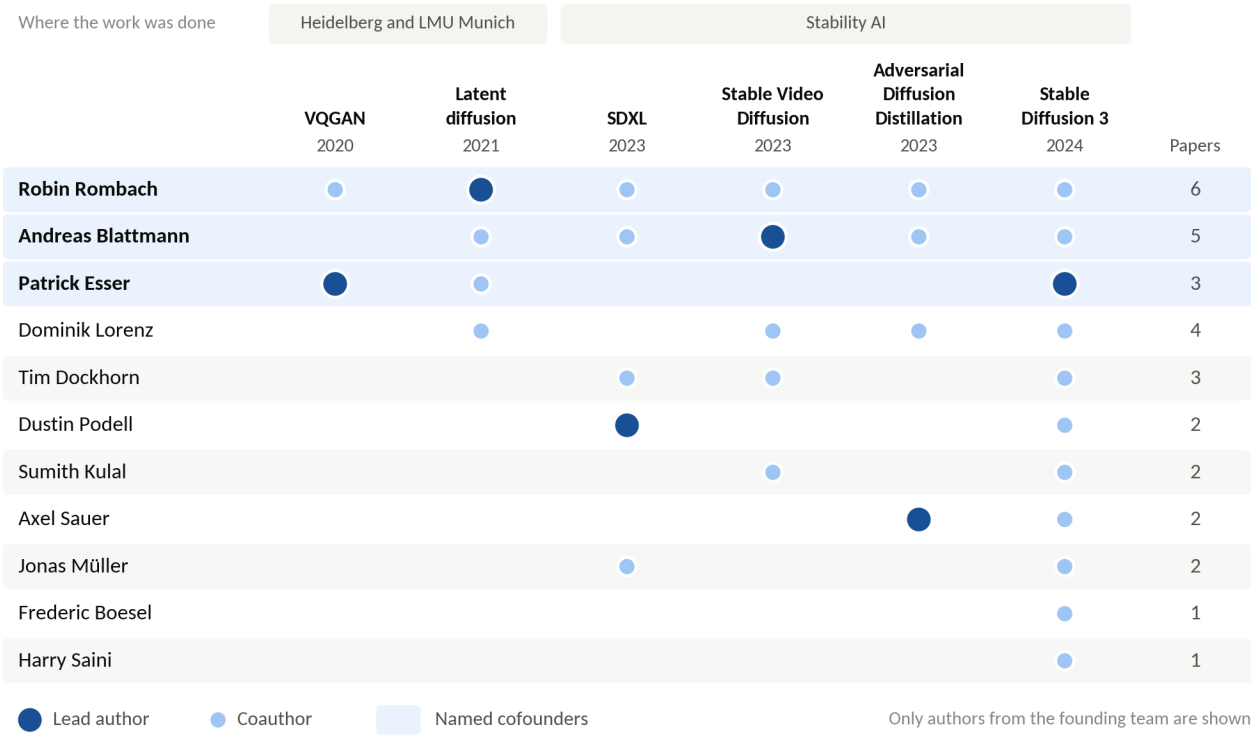


Figure 2: Authorship of six landmark papers by the founding team. Sources: arXiv (VQGAN, latent diffusion, SDXL, Stable Video Diffusion, ADD, Stable Diffusion 3)

By early 2024, Stability AI was in crisis. It had lost several senior leaders, and internal documents reportedly showed expenses far above revenue. In March 2024, Rombach, Blattmann and Lorenz left the company, and days later CEO Emad Mostaque resigned. Rather than join another lab, the team started its own company in Freiburg, a university town on the edge of the Black Forest and far from the usual tech hubs. “I am deeply convinced that you don’t need to only be in a major tech hub to execute and build a company,” Rombach told Germany Trade & Invest, calling the choice of Freiburg “the best thing that we have done so far.”

Black Forest Labs came out of stealth on August 1, 2024 with FLUX.1, a family of three image models, and a $31 million seed round led by Andreessen Horowitz. Angel investors included former Oculus CEO Brendan Iribe, Y Combinator CEO Garry Tan, Hollywood dealmaker Michael Ovitz and AI researchers Timo Aila and Vladlen Koltun. Within two weeks, FLUX was powering image generation in xAI’s Grok, and independent rankings at the time put its quality ahead of Midjourney and OpenAI. In December 2025, the company raised a $300 million Series B at a $3.25 billion valuation, led jointly by Salesforce Ventures and Anjney Midha’s firm AMP, with Temasek, Canva and Figma Ventures among the new investors. That brought total funding to about $450 million.

From a university lab to FLUX 3

Key research, product and business milestones, 2020 to 2026

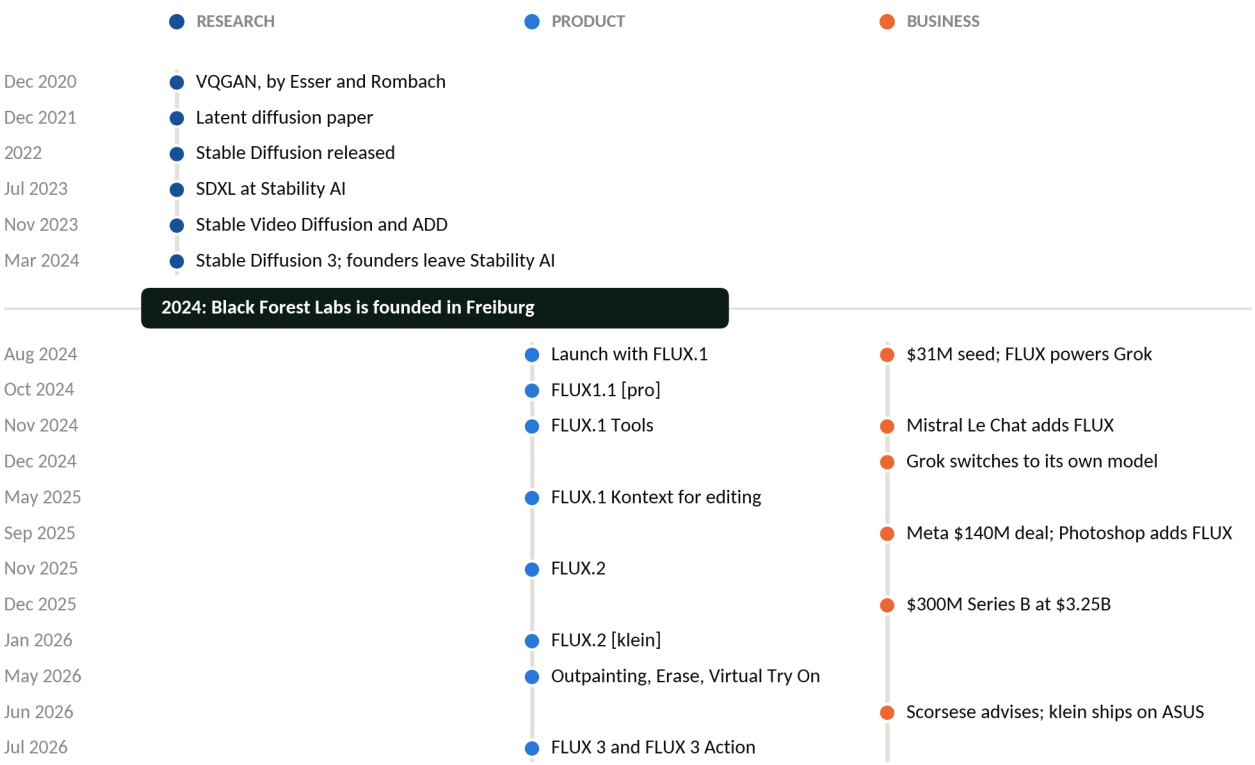


Figure 3: Milestones from the research years to FLUX 3. Sources: [arXiv](), [bfl.ai](), [Wikipedia](), [TechCrunch](), [Bloomberg]()

About $450M raised in 16 months, with partners on the cap table

Capital by round (left) and investors by type (right)

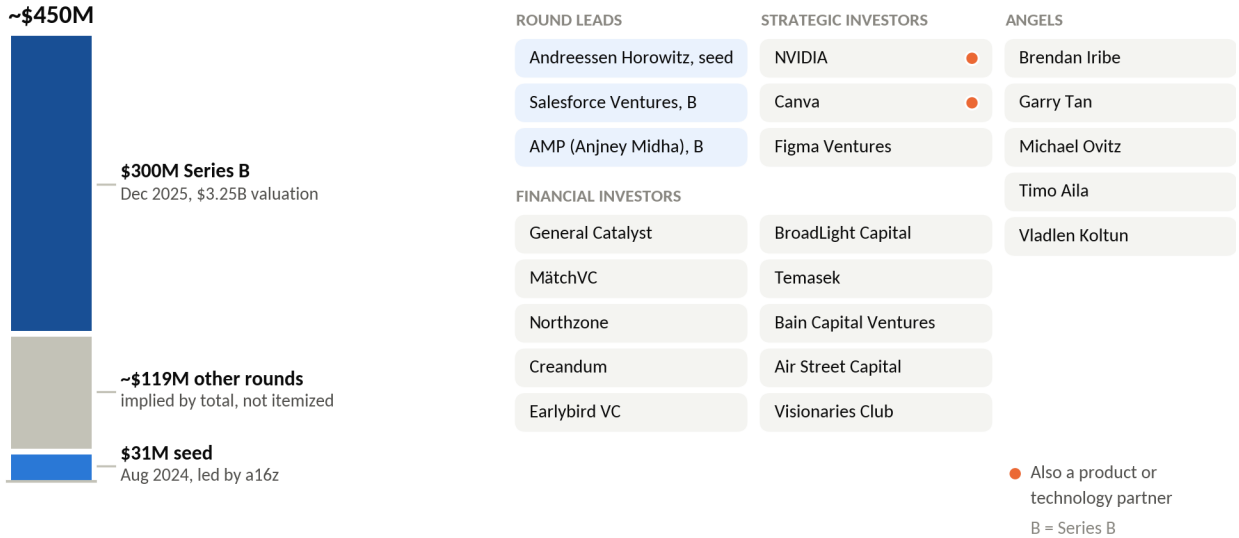


Figure 4: Funding by round and investors by type; the middle segment is implied by the reported total. Sources: [bfl.ai seed announcement](), [bfl.ai Series B](), [Crunchbase News]()

What sets the team apart is that it has repeatedly produced the core techniques of modern image generation instead of adapting other people's: VQGAN, latent diffusion, SDXL, Stable Video Diffusion, Adversarial Diffusion Distillation and the rectified flow transformer behind Stable Diffusion 3 all came out of this group. Each founder also brings a distinct specialty. Rombach was lead author on latent diffusion and now runs the company; Esser was lead author on Stable Diffusion 3 after leading applied research at Runway; Blattmann was lead author on Stable Video Diffusion; Podell was lead author on SDXL; and Sauer's distillation research made fast generation practical. The result is an unusually fast shipping cadence for a team the company itself calls "small but exceptional": FLUX.1, FLUX.2 and the multimodal FLUX 3 all launched within two years.

The company has also drawn unusual talent into its orbit. Its first advisory board included Ovitz and Tübingen professor Matthias Bethge. Anjney Midha joined the board as a general partner at Andreessen Horowitz and later led the Series B together with Salesforce Ventures through his own firm. In June 2026, Martin Scorsese became an advisor after using FLUX to storyboard scenes with his team. "Now, with this tool, I can share what I'm visualizing more clearly and efficiently to my creative team," Scorsese said.

The founding team's technical talent comes from a single lineage: Ommer's academic group and the generative research team at Stability AI. As the business matured, Black Forest Labs added a commercial layer around that research core rather than replacing it. It now runs solutions engineers and forward deployed engineers who embed directly with customers, for example building transparent background endpoints for Microsoft and running human evaluations on use cases such as Amazon product photography.

How Black Forest Labs is set up

Legal entities, leadership and the path from research to customers

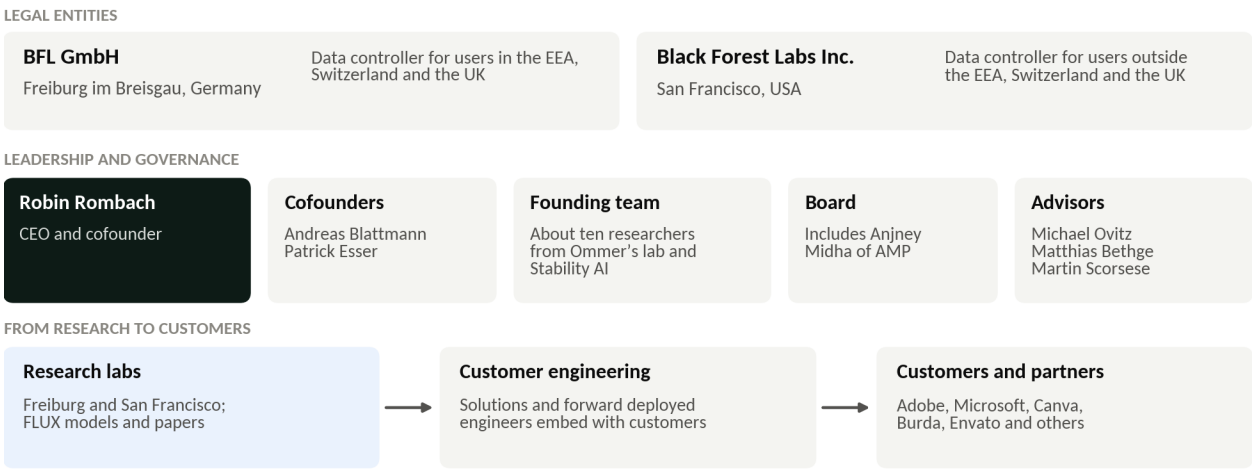


Figure 5: Company structure. Sources: bfl.ai privacy policy, bfl.ai, a16z, GTAI, bfl.ai blog

There have been no publicly reported changes to the CEO, CTO, COO or President roles since founding. Rombach still leads the company as CEO, and at the FLUX 3 launch in July 2026 he framed its

next chapter this way: “True intelligence means perceiving the world: predicting how it will change, taking action, and learning from results.”

Product

Black Forest Labs sells what it calls visual intelligence: AI models that create and edit images, and increasingly video, sound and even robot movements. The offering is easiest to understand in three layers: the FLUX model families, a set of single purpose tools built on top of them, and three ways to buy access (a pay as you go API, open weights under paid licenses, and enterprise deployments). Many people who use FLUX never see the Black Forest Labs name, because they meet it inside other products such as Adobe Photoshop, Mistral’s Le Chat, Canva or Picsart.

Table 1: The Black Forest Labs product suite at a glance

Layer	What it does	Main offerings
FLUX.2 (November 2025)	Creates and edits still images for production use	[max], [pro], [flex], [dev] with open weights, and the small, fast [klein]
FLUX 3 (July 2026)	One model for video with sound, images and robot actions	FLUX 3 Video, FLUX 3 Image, FLUX 3 Action, FLUX 3 Dev
FLUX Tools (2026)	Single purpose edits, each in one API call	Video Edit, Video Upscaler, Outpainting, Erase, Virtual Try On
Access	How customers pay for and run the models	API and Playground, commercial licenses for open weights, enterprise deployments

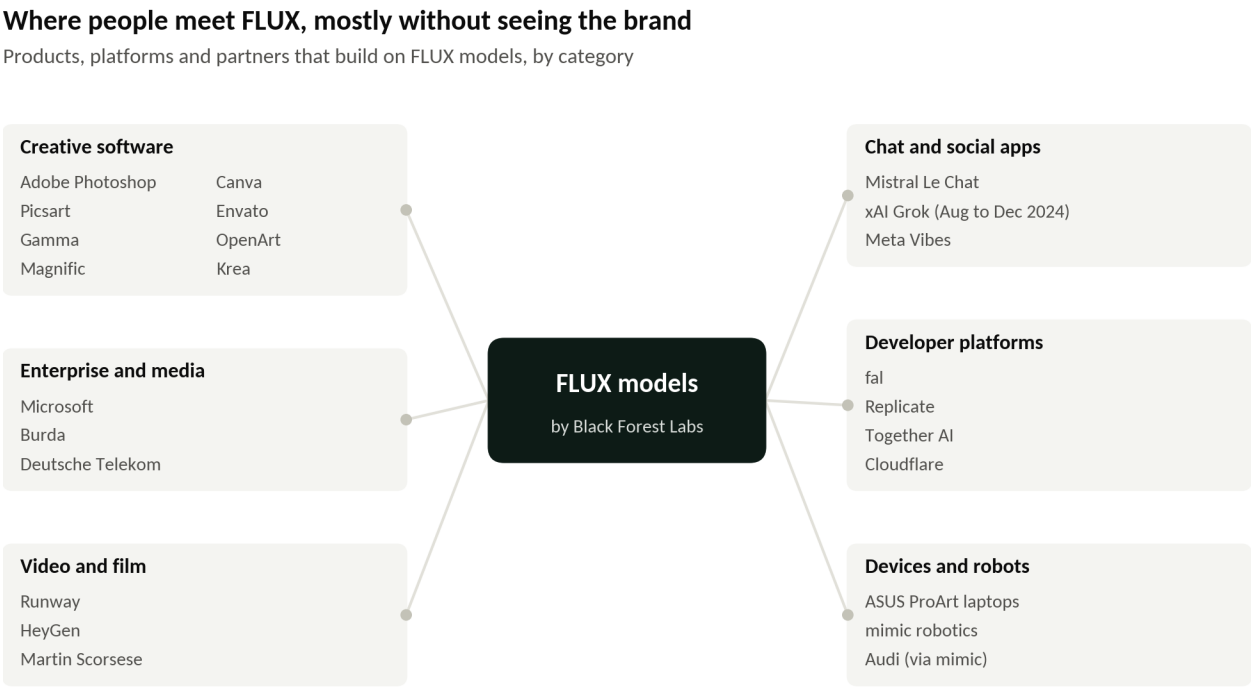


Figure 6: Where FLUX shows up. Sources: bfl.ai, bfl.ai FLUX 3, Wikipedia, GTAI

FLUX.2: images for production

Launched on November 25, 2025, FLUX.2 is the company's workhorse for still images. A single model both creates images from a text prompt and edits existing ones, at up to four megapixels in any aspect ratio. Its standout feature is control, which is what professional users care about most:

- **References instead of retraining.** Users can combine up to ten reference images to keep a character, product or style consistent across a whole campaign, without training a custom model.
- **Brand accuracy.** Colors can be specified as exact hex codes, and the model renders readable typography, which makes it usable for ads, packaging mockups and interface designs.
- **Structured prompts.** Besides plain language, prompts can be written in a structured JSON format and combined with pose guidance, which helps teams automate large batches of images.

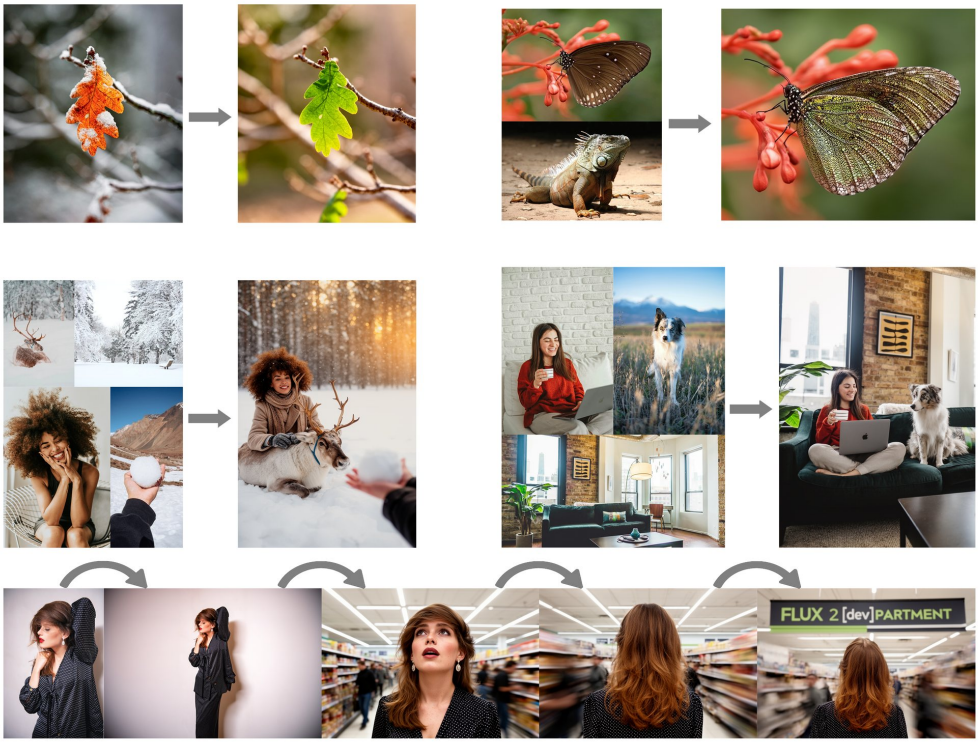


Figure 7: FLUX.2 editing. Top: changing a single detail, and merging elements of two photos. Middle: combining several reference photos into one scene. Bottom: a chain of edits that keeps the same person consistent. Source: Black Forest Labs

Under the hood, FLUX models are rectified flow transformers. In plain terms, the model starts from random noise and learns the most direct path to an image that matches the prompt, and it does this on a compressed version of the image, the latent diffusion idea the founders introduced. A more direct path means fewer steps, and fewer steps mean faster and cheaper generation.

FLUX.2 comes in five variants that trade quality against speed and openness: [max] for the highest quality, including "grounded" generation that pulls in real time information from the web; [pro] as the production default; [flex] for typography and fine detail; [dev], a 32 billion parameter model with open weights for noncommercial use; and [klein].

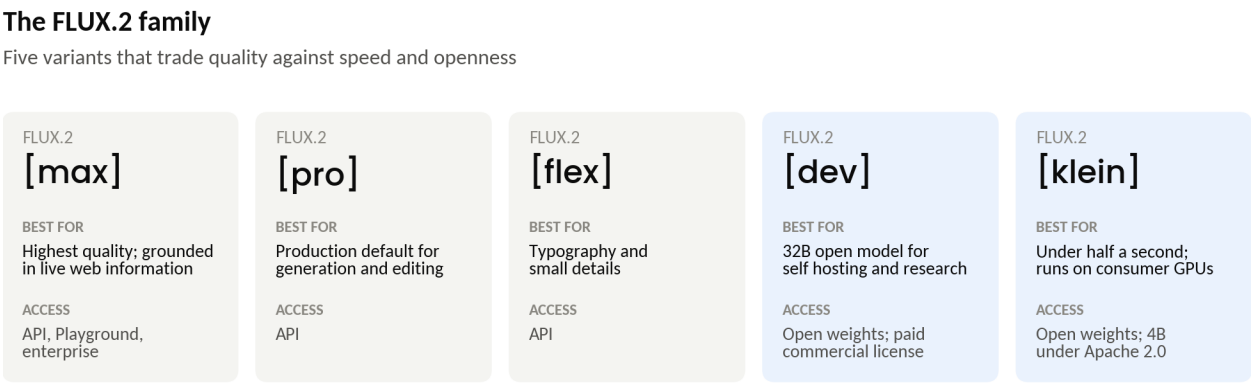


Figure 8: The FLUX.2 variants. Sources: bfl.ai FLUX.2, bfl.ai FLUX.2 [max], bfl.ai licensing

FLUX.2 [klein], released in January 2026, is the most strategically interesting variant. Its 4 billion and 9 billion parameter models generate or edit an image in under half a second on modern hardware and run on consumer graphics cards, and the 4 billion parameter version is released under the permissive Apache 2.0 license. In June 2026, klein became the first FLUX model to ship on consumer hardware, running directly on ASUS ProArt laptops without any connection to an API. Speed at this level turns image generation from a "submit and wait" task into something interactive, closer to sketching.

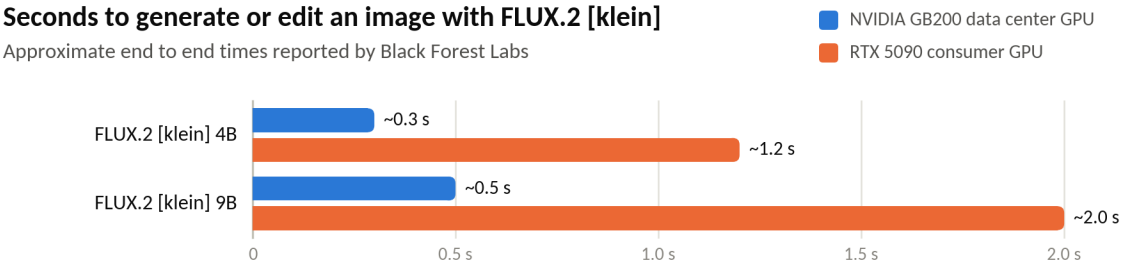


Figure 9: FLUX.2 [klein] speed on data center and consumer GPUs. Source: bfl.ai



Figure 10: Images generated with FLUX.2 [klein]. Source: Black Forest Labs

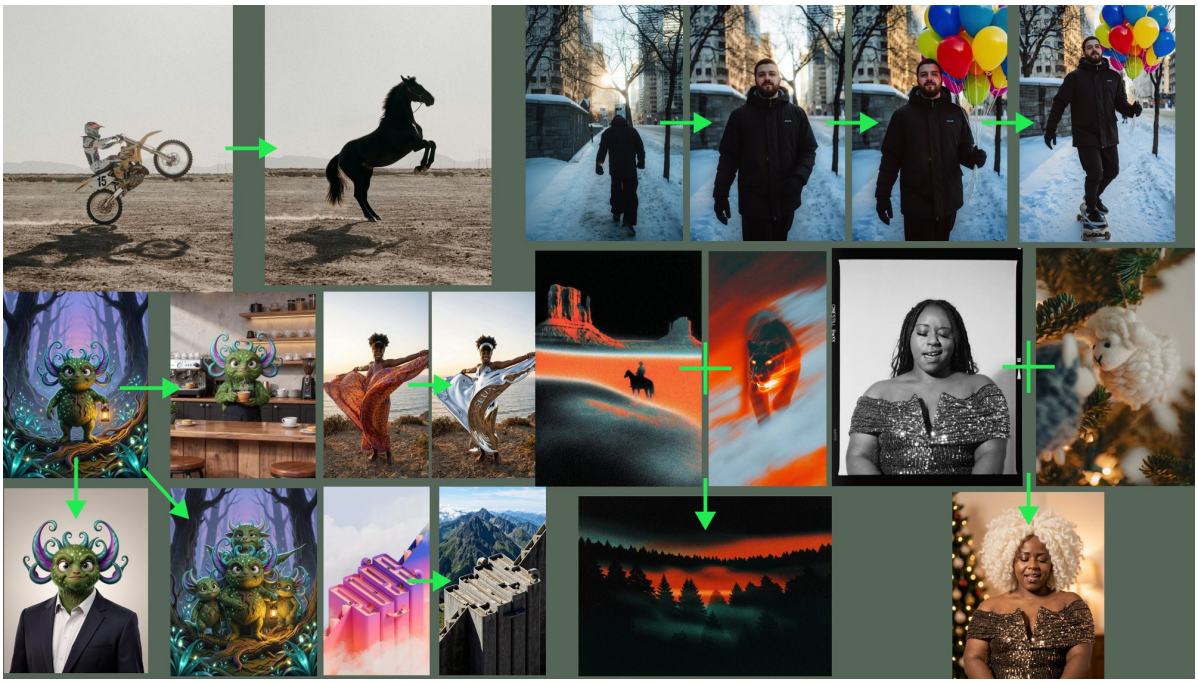


Figure 11: Edits made with FLUX.2 [klein]: replacing objects, chaining several edits on the same person, and combining elements of two images. Source: Black Forest Labs

FLUX 3: one model for video, sound, images and robots

FLUX 3, launched in early access on July 23, 2026, is the company's bet on a single model for every medium. Instead of training separate models for images, video and audio, FLUX 3 learns from all three together, which the company argues gives it a better internal picture of how the world looks, sounds and moves. It comes in four parts:

- **FLUX 3 Video** creates clips of up to 20 seconds in up to 4K from text, images or existing video, and can continue a clip or follow set keyframes. Sound is generated natively at no extra cost, including speech with accurate lip sync and multilingual dialogue. Pricing is pay as you go, from $0.06 per second of draft HD video to $0.80 per second in 4K. In the company's own preliminary tests, reviewers preferred it over the latest models from Runway and Luma in 77% and 93% of comparisons.
- **FLUX 3 Image** is the next generation of the image model, rolling out after video.
- **FLUX 3 Action** is a 7 billion parameter model with open weights that looks at camera images and predicts both what happens next and which motor commands a robot should execute. It is rolling out through partners, starting with mimic robotics.
- **FLUX 3 Dev** is an open weights version of the shared backbone that developers will be able to run and adapt themselves.

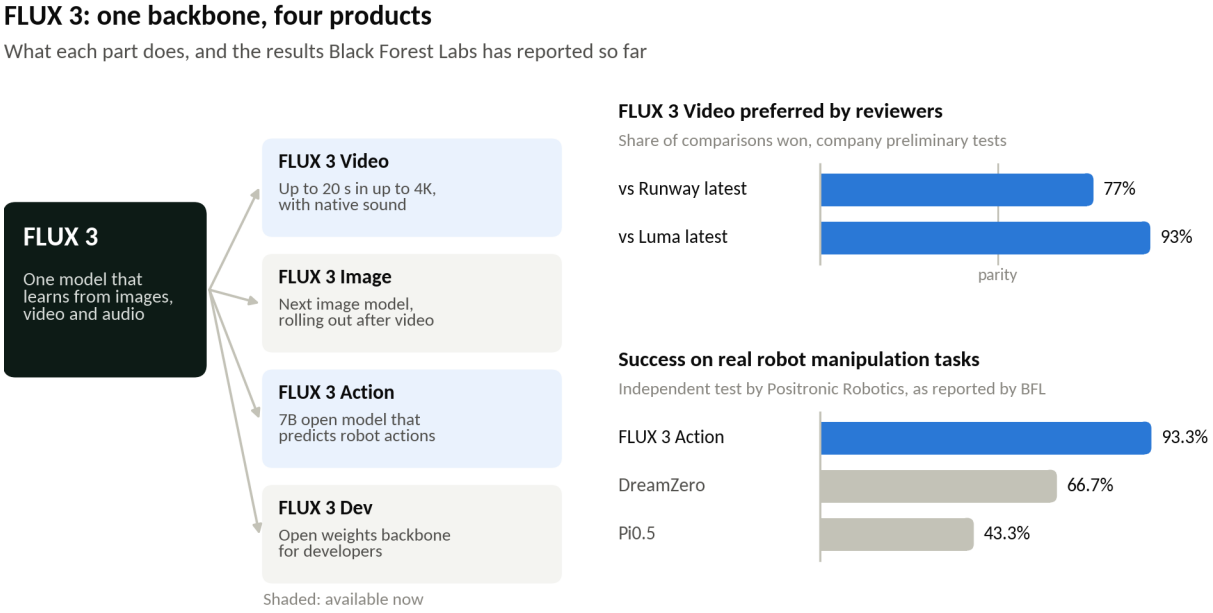


Figure 12: FLUX 3 components and company reported results. Sources: bfl.ai FLUX 3, bfl.ai FLUX 3 Action, VentureBeat

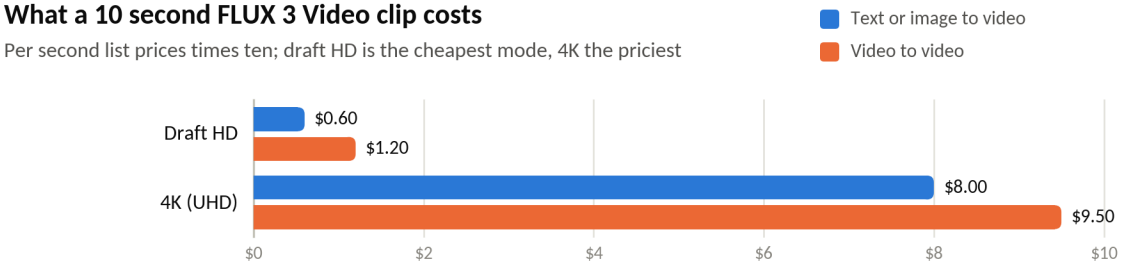


Figure 13: FLUX 3 Video list prices for a 10 second clip. Source: bfl.ai

FLUX Tools: one API call per job

Alongside the models, Black Forest Labs released a set of single purpose tools in 2026, each available as one API call. Video Edit changes objects, settings, style, text or dialogue in an existing clip with a prompt; the Video Upscaler lifts footage to 2K or 4K; Outpainting extends an image past its frame without a prompt; Erase removes objects together with their shadows and reflections; and Virtual Try On dresses any person in any garment with logos and prints intact, in under four seconds. Tools like these package the raw model into finished outcomes that product teams can plug in without any prompt engineering.

How customers access FLUX

Black Forest Labs sells the same models through three channels:

- **API and Playground.** Developers pay as they go, with no subscriptions or seat fees, and can test models for free in a web Playground. FLUX is also sold through inference platforms such as fal, Replicate and Together AI, where the company says its models are among the most widely adopted.

- **Open weights with licenses.** Several models can be downloaded from Hugging Face, where FLUX.1 [dev] alone has about 15,000 likes and the company says its open models are the most popular image models on the platform. Research and personal use is free under noncommercial terms; companies that want to run the weights inside their own products buy a self serve commercial license (Table 2).
- **Enterprise.** Large customers can use a managed API with zero data retention, host open weights on their own servers for full control over their data, or develop custom models together with Black Forest Labs on dedicated, isolated instances, including a white labeled studio interface. The offering covers SOC 2 Type II, ISO 27001 and GDPR requirements.

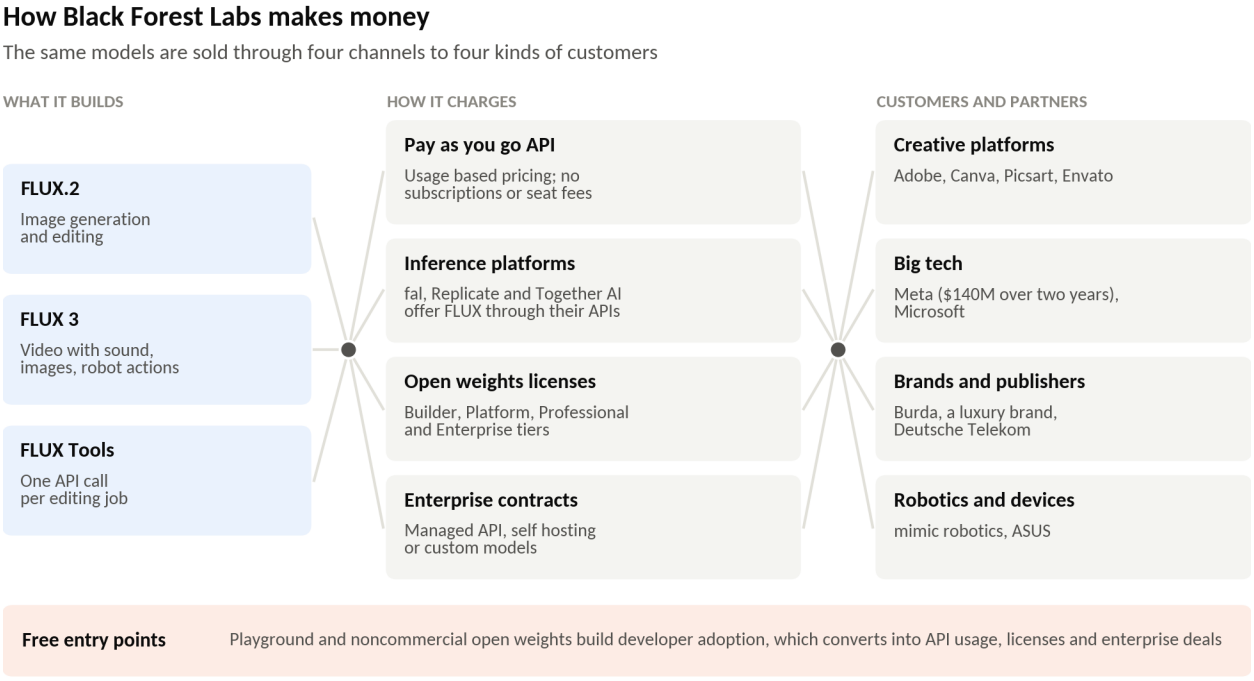


Figure 14: Business model. Sources: bfl.ai pricing, bfl.ai licensing, bfl.ai enterprise, bfl.ai Series B, Bloomberg

Table 2: Self serve commercial licenses for open weights (Source: Black Forest Labs)

Tier	Built for	Images per month	Scope
Builder	Developers and early stage teams	10,000	1 domain, 10 users, fine tuning and LoRA rights
Platform	SaaS companies and product teams	100,000	1 domain, 10 users, fine tuning rights
Professional	Creative agencies and studios	100,000	Up to 3 domains and the first 3 clients included
Enterprise	Organizations with bespoke needs	Custom	All models and new releases, custom terms

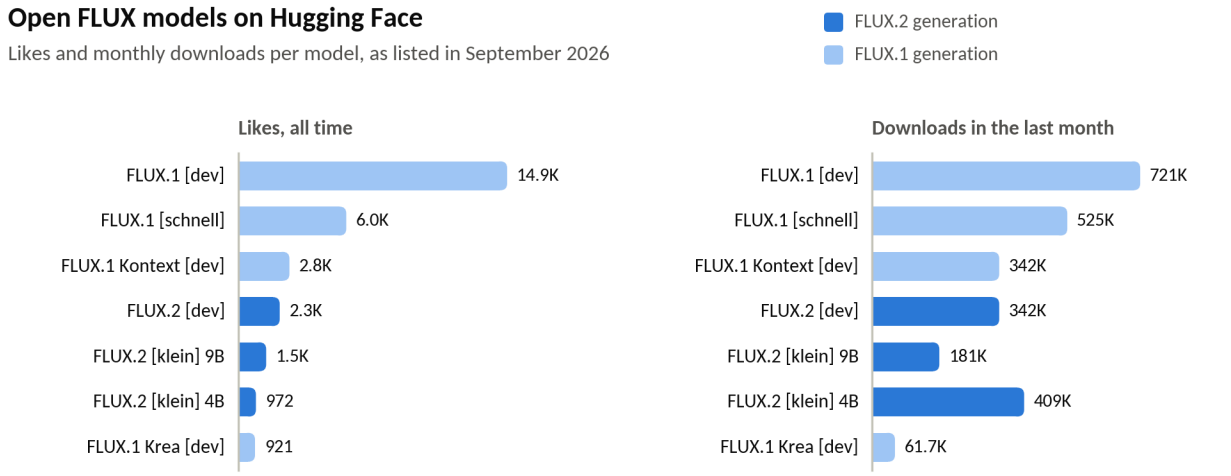


Figure 15: Adoption of open FLUX models on Hugging Face. Source: Hugging Face

Value proposition

For customers, the pitch combines three things that are rarely available together: frontier quality, fine control and freedom over where the model runs. The same FLUX family can be called through an API, run on a company's own servers or even on a laptop, and each option comes with clear commercial terms. For a platform like Adobe or Canva, that means state of the art features without depending on a single closed provider. For an enterprise, it means brand safe content on its own infrastructure.

Job to be done

The job differs by customer, but it always ends in production rather than experimentation:

- **Creative platforms** such as Adobe, Canva, Picsart and Envato: *give my users best in class image and video generation inside my product, without training a frontier model myself.*
- **Brands and publishers** such as Burda or a luxury brand: *produce far more on brand visuals, far faster, without losing consistency.*
- **Robotics teams** such as mimic robotics: *teach robots new tasks from video instead of months of demonstrations.* As mimic's CTO put it, a model built on video "picks up new tasks in minutes, not days."

Aspirin or vitamin?

For creative software platforms, FLUX has become an aspirin. Users now expect generative fill, clean cutouts and endless on brand variations, and Adobe adding FLUX.1 Kontext [pro] to Photoshop's generative fill shows how quickly these features became table stakes. For enterprises producing marketing content, FLUX is moving from vitamin to aspirin as volumes grow: one luxury brand customer went from 1,000 to more than 100,000 images per season, cutting the time per asset from two to four hours of manual work to under a minute. The risk sits on the platform side. A model that powers someone else's product can be swapped out, as Grok showed when it replaced FLUX with xAI's own Aurora model in December 2024.

Real world examples

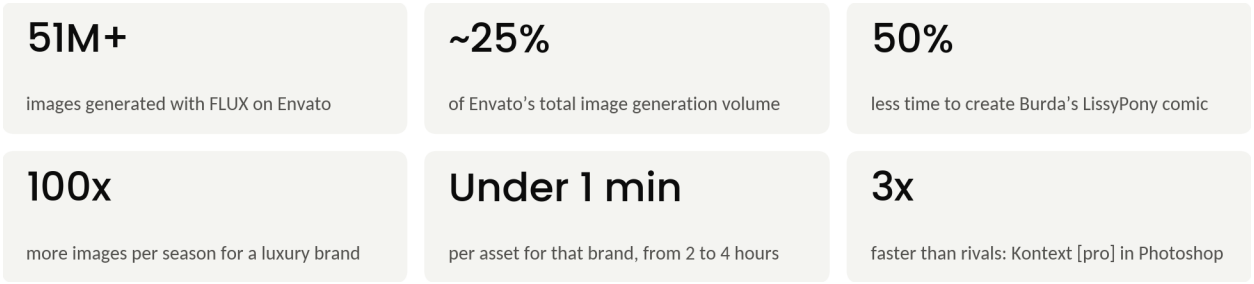


Figure 16: Results reported by Black Forest Labs and its customers. Sources: bfl.ai Envato, bfl.ai Burda, bfl.ai enterprise, bfl.ai Photoshop

- **Adobe.** Photoshop beta users can choose FLUX.1 Kontext [pro] inside generative fill, which Black Forest Labs says runs three times faster than competing models.
- **Envato.** The creative asset marketplace has generated more than 51 million images with FLUX, about a quarter of its total image generation volume, and delivers five variations of a stock image in under eight seconds.
- **Burda.** The German publisher used finetuned FLUX.1 models together with Depth, Kontext and Fill to produce its children's comic LissyPony, cutting creation time by 50% for its first fully AI supported issue in September 2025.
- **Microsoft.** Forward deployed engineers built transparent background endpoints for Microsoft, one of the partners the team works with directly.
- **Martin Scorsese.** Used FLUX to storyboard scenes before shooting, saying it let his team move faster "without sacrificing quality or craft."
- **mimic robotics.** Built a joint video and action model on FLUX 3 that is being used for robot manipulation tasks at Audi.
- **ASUS.** Ships FLUX.2 [klein] directly on ProArt laptops, the first FLUX model on consumer hardware.