Atlas Cloud

ATLAS CLOUD RESEARCH

# The Model Mix 2026

Which image, video, audio and language models builders use, how they use them, and where the market goes in 2027. Usage data from Atlas Cloud, January to October 2026.

By Colin Mo, Head of Product Marketing, Atlas Cloud

First edition, November 2026

**Contents**

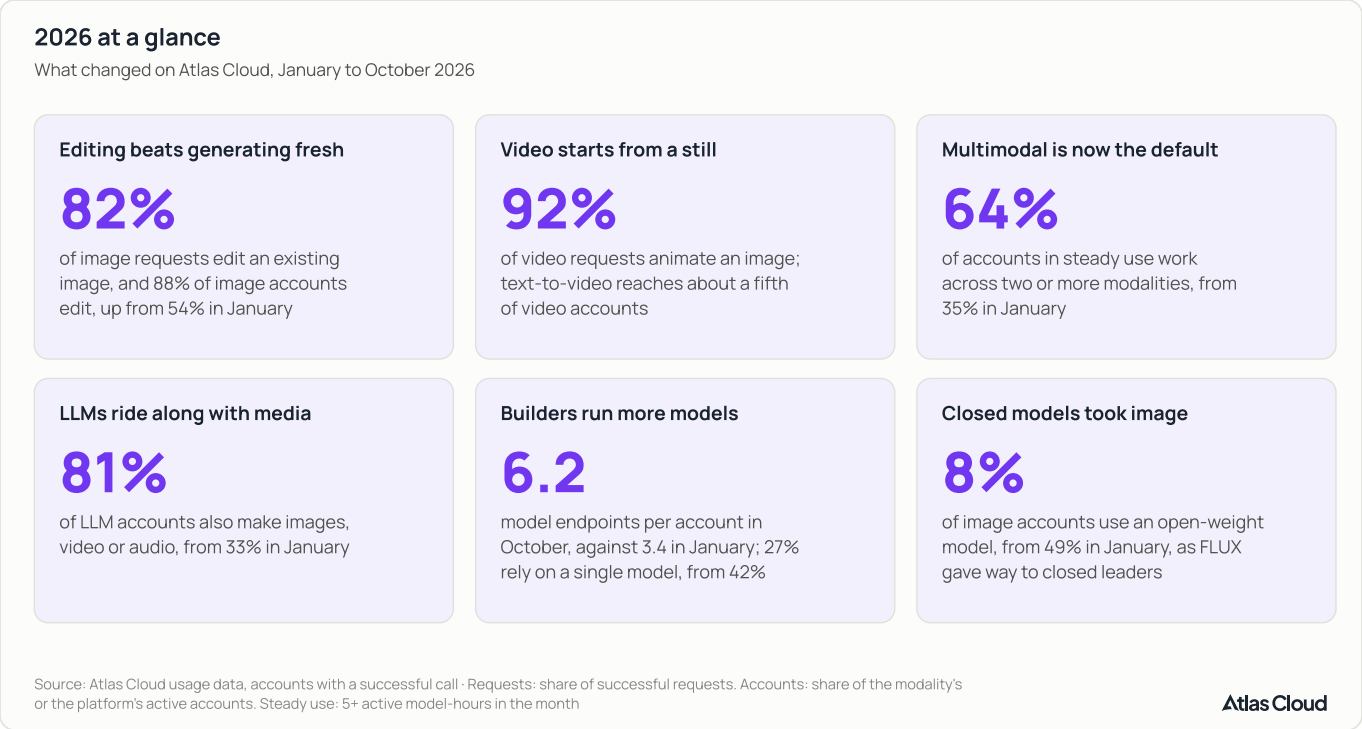
Summary
About this report
Video
Image
Audio
3D
Language models
What it means by industry
What it means
Sources
About the author
About Atlas Cloud

# Summary

Most generative media work on Atlas Cloud now starts from an image, and the people doing it changed in 2026: companies moved from testing to steady use, and individual builders more than doubled.

- **88%** of image accounts used editing in October, up from 54% in January, and **96%** of video accounts animated an image.
  - *Our reading:* a starting image pins down the product, the brand and the composition in a way a prompt alone does not, so teams get usable output in fewer tries.
- **Companies moved from testing to steady use.** In January, 18% of business accounts were in steady use; in October, **67%** were. They run more models than other accounts (10.5 endpoints against 7.6) and keep about one in four of the models they try.
  - *Our reading:* brands and e-commerce teams put product images in and get versions out, studios and short-drama producers drive video volume, and software teams start on language models and add media.
- **Business use deepened fastest from April.** The number of business accounts in steady use grew about **five and a half times** from January to October, fastest from April to July, alongside new video models: Dreamina Seedance 2.0 and Wan 2.7 in April, then Dreamina Seedance 2.5 and Wan 3.0 in August, which now sit near the top of the video boards.
  - *Our reading, from the studios and brands we work with:* better video models likely helped, as filmmaking, studio, advertising and other media-based work grows now that models clear the bar those teams need. The usage data alone can't show the cause.
- **Individual builders grew faster than companies.** Personal-email accounts **more than doubled**, most of it between May and August; business accounts grew about half.
  - *Our reading:* maturing tools and falling prices put capable models within reach of creators, prosumers and developers, who experiment on their own projects, some of them building influencer accounts.
- **Builders work across media:** 64% of accounts in steady use touched two or more modalities in October, up from 35% in January, 81% of language-model accounts also make media, and audio went from almost nothing to **one account in ten**.
- **The lead changed hands** in image, video, text-to-speech and 3D, and builders add the new leader before they drop the old one.
- **Usage doesn't simply follow benchmarks or price:** on none of six leaderboards was the top-rated model we could match the most used, and price barely related to share, apart from one heavy cut.

FIGURE K.1




## Where the market goes in 2027

Each prediction starts from a trend we measured across the year; the business effects are our judgment.

### More production starts from an existing image

- Editing reached 88% of image accounts in October, from 54% in January, and image-to-video 96% of video accounts, from 84%. Both climbed through the spring and have held near those highs since.
- We expect the trend to continue as more brand and product work moves to generated media, with marketing and e-commerce teams that hold large libraries of approved images moving fastest.

### Video becomes everyday work for brands

- Video was used by **72%** of active accounts in October, against 58% in January.
- Business accounts moved toward steady use: **67%** were in steady use across modalities in October, against 18% in January.
- This tracks with how enterprises are exploring AI content for brands, ads, commercials and products, with studios starting to experiment with film (*our reading*).

### Language models get bought as part of media work

- **81%** of LLM accounts on Atlas also make media, up from 33% in January.
- The median such account made **28 text calls** in October: the LLM is a helper, likely for scripts, captions, ad copy or prompts (*our inference*).
- We expect more LLM use from media teams as this continues.

### Builders will keep several models in use

- The leader changed in **image, video, text-to-speech and 3D** this year.
- The average account used 6.2 model endpoints in October against 3.4 in January, and the share relying on a single model fell from 42% to **27%**.
- We expect buyers to build so they can change models, and to test each new one on their own work. *Atlas sells multi-model access and has an interest in this one.*

### 3D builders go deeper before the base widens

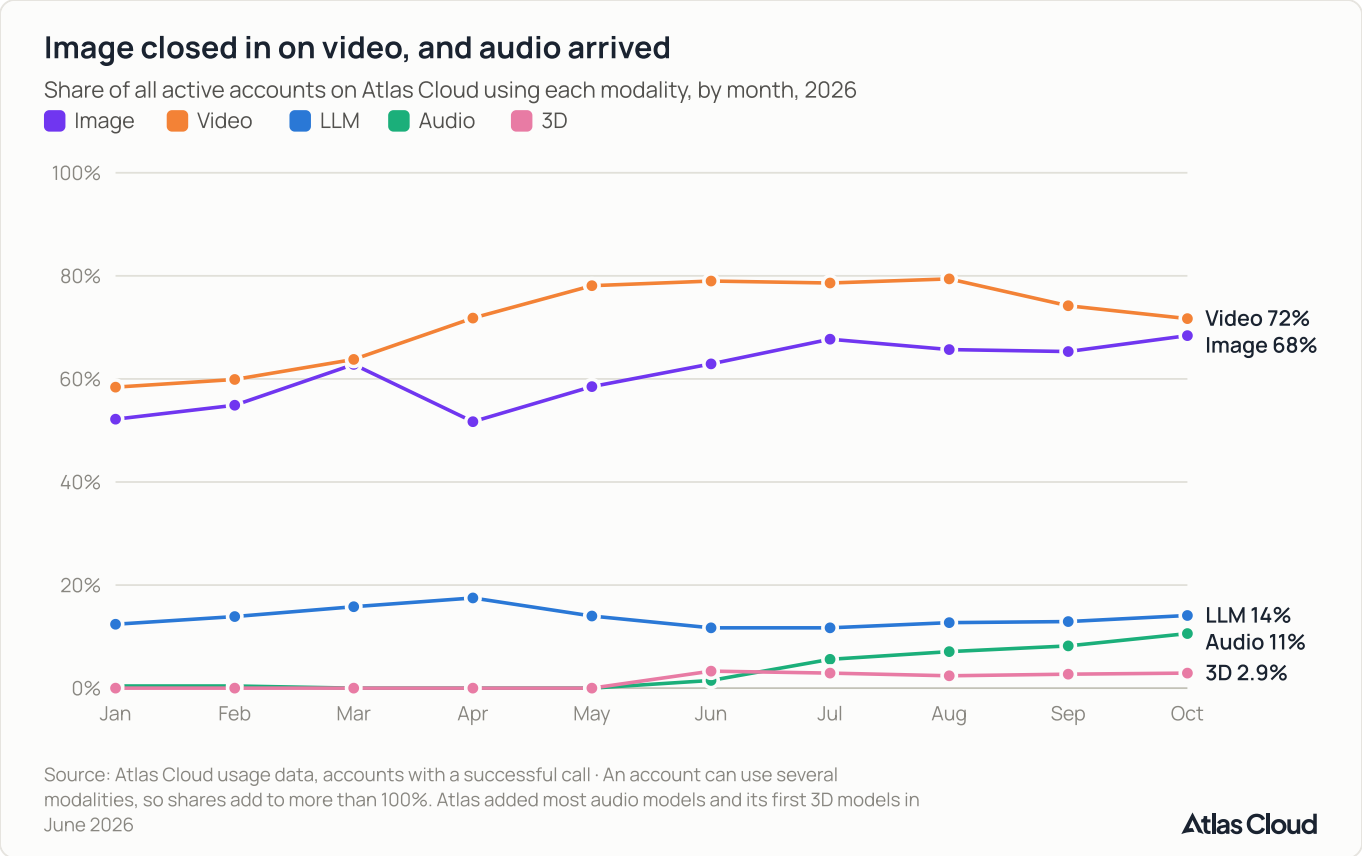
- 3D has been about **3%** of active accounts since Atlas added it in June, but the average 3D account made about **4.7 times** as many successful requests in October as in June.
- 38% of 3D accounts used two or more 3D families in October, up from 15% in June. We expect this deeper use to continue before the base widens, since a 3D asset can be reused from any angle where a video clip is a fixed series of frames (*our reading*). We are watching 3D rather than calling it.
- *Atlas's 3D lineup is thin today (see 3D), so our numbers may understate 3D demand.*

## The year in numbers

### Image closed in on video, and audio arrived

- **Video and image:** in January, 52% of active accounts on Atlas used an image model and 58% a video model. Video held at about 79% from May to August, then eased to 72% in October after Atlas retired some Wan endpoints in September; image rose to **68%** as editing spread.
- **Audio** went from under 1% to **11%** of accounts after Atlas added most audio models in June. Builders added it to image and video work rather than buying it on its own (see Audio).
- **Language models** stayed at 12 to 18% of accounts, mostly as a helper for media work.
- **3D** stayed at about 3%, and about one in five of Tripo H3.1's first users used 3D again a month later.

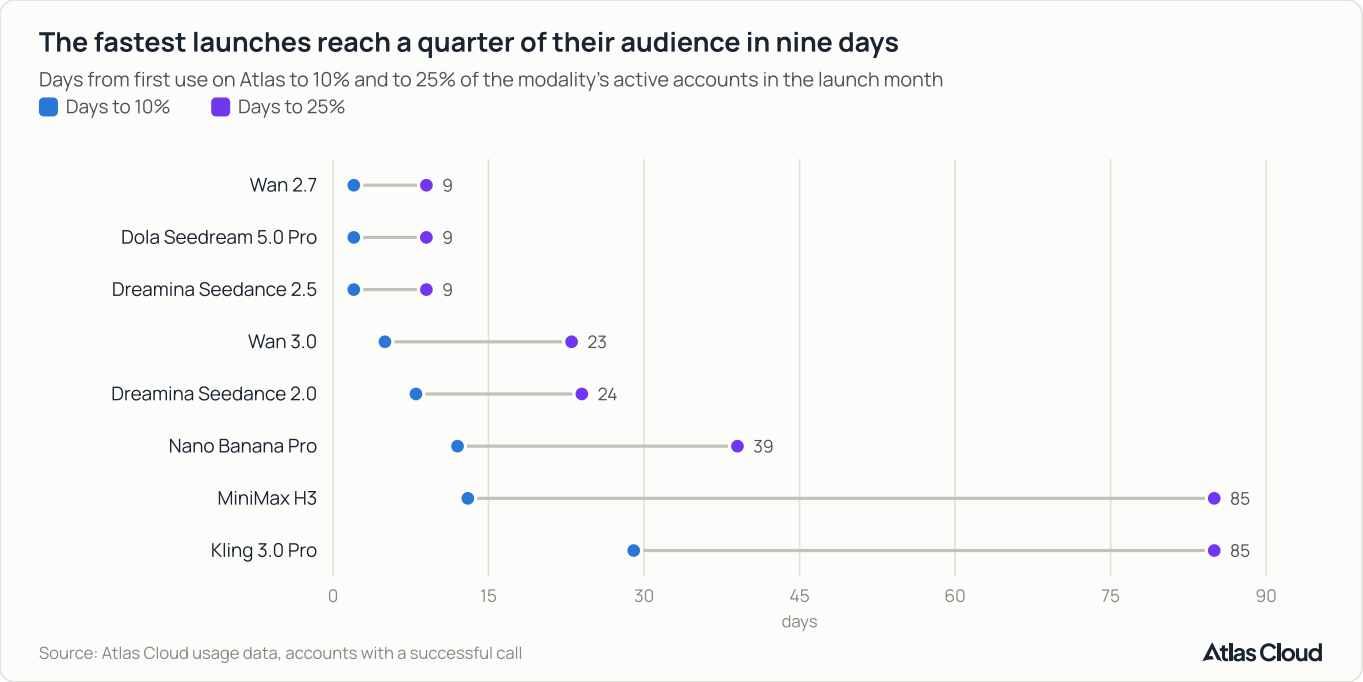
FIGURE M.1




## New models get picked up in days

- The fastest image and video launches since late 2025 reached a quarter of their modality's launch-month accounts within **nine days** of first use on Atlas. Dreamina Seedance 2.0 took 24 days, Nano Banana Pro 39 and Kling 3.0 Pro 85.
- The leading launches brought in new builders more than they pulled them off rivals. Of the accounts that made Dreamina Seedance 2.5 their main video model in its first two months, **52%** had not made video on Atlas before and 23% moved up from an older Dreamina Seedance. For Dola Seedream 5.0 Pro, 63% were new to image on Atlas.
- *Atlas promotes its launches, so this is how Atlas's audience responds.*

FIGURE G.2




## What changed: who leads

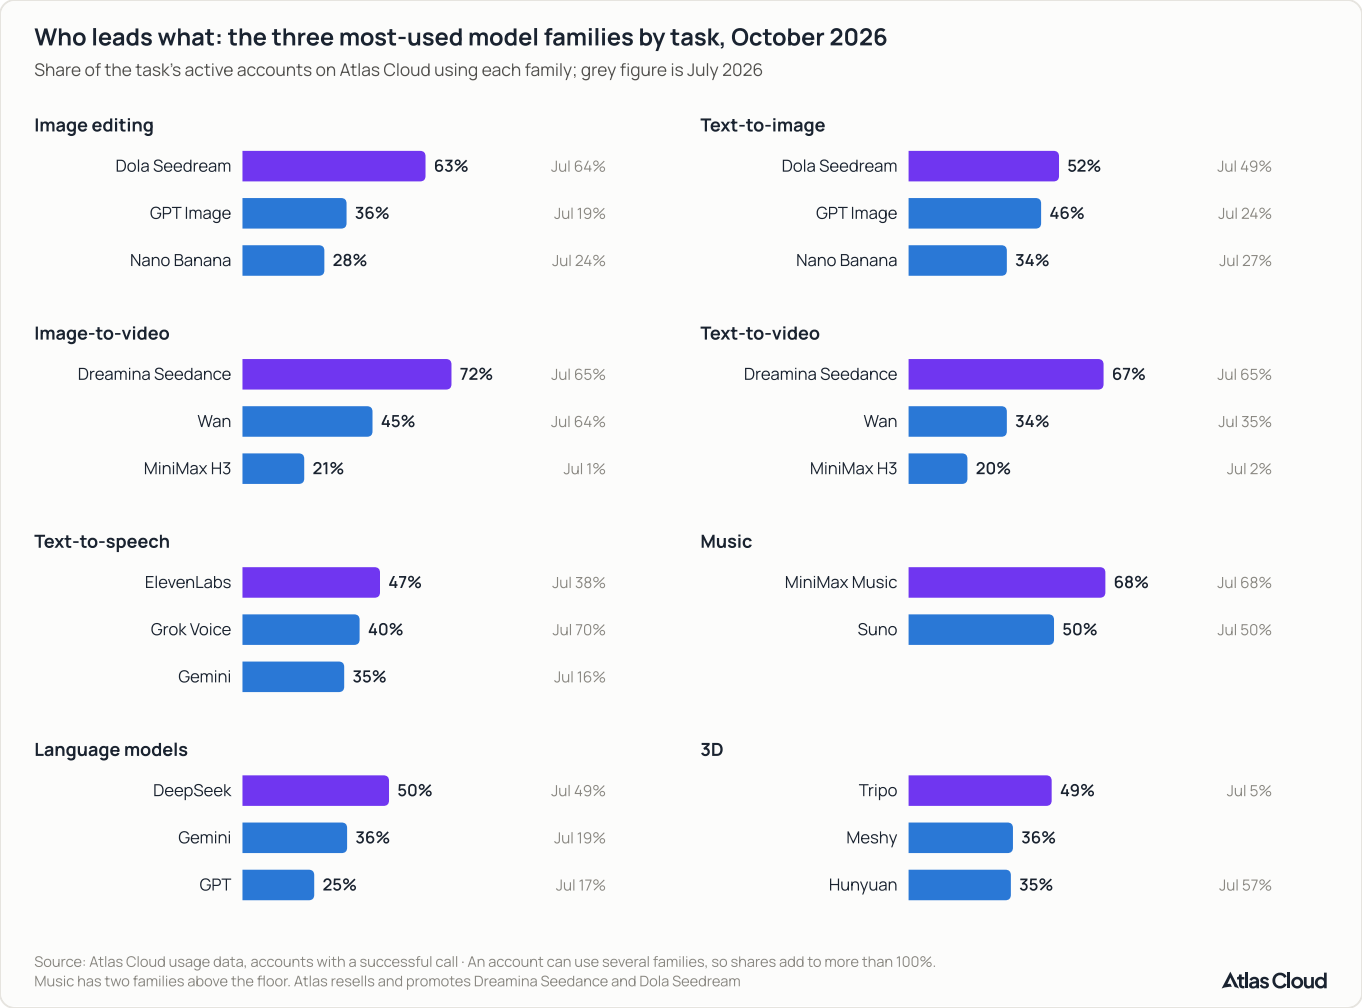
**The lead changes hands, and builders add the new leader before they drop the old one.**

**What it means (*our reading*):** a model maker keeps the lead only while its newest version stays ahead, and builders can move quickly because the old leader usually stays available.

### One family leads each media task

In October, Dola Seedream led image editing and text-to-image, Dreamina Seedance led image-to-video and text-to-video, ElevenLabs led text-to-speech, MiniMax Music led music, DeepSeek led language models and Tripo led 3D.

FIGURE T.1




## The lead moved in most tasks

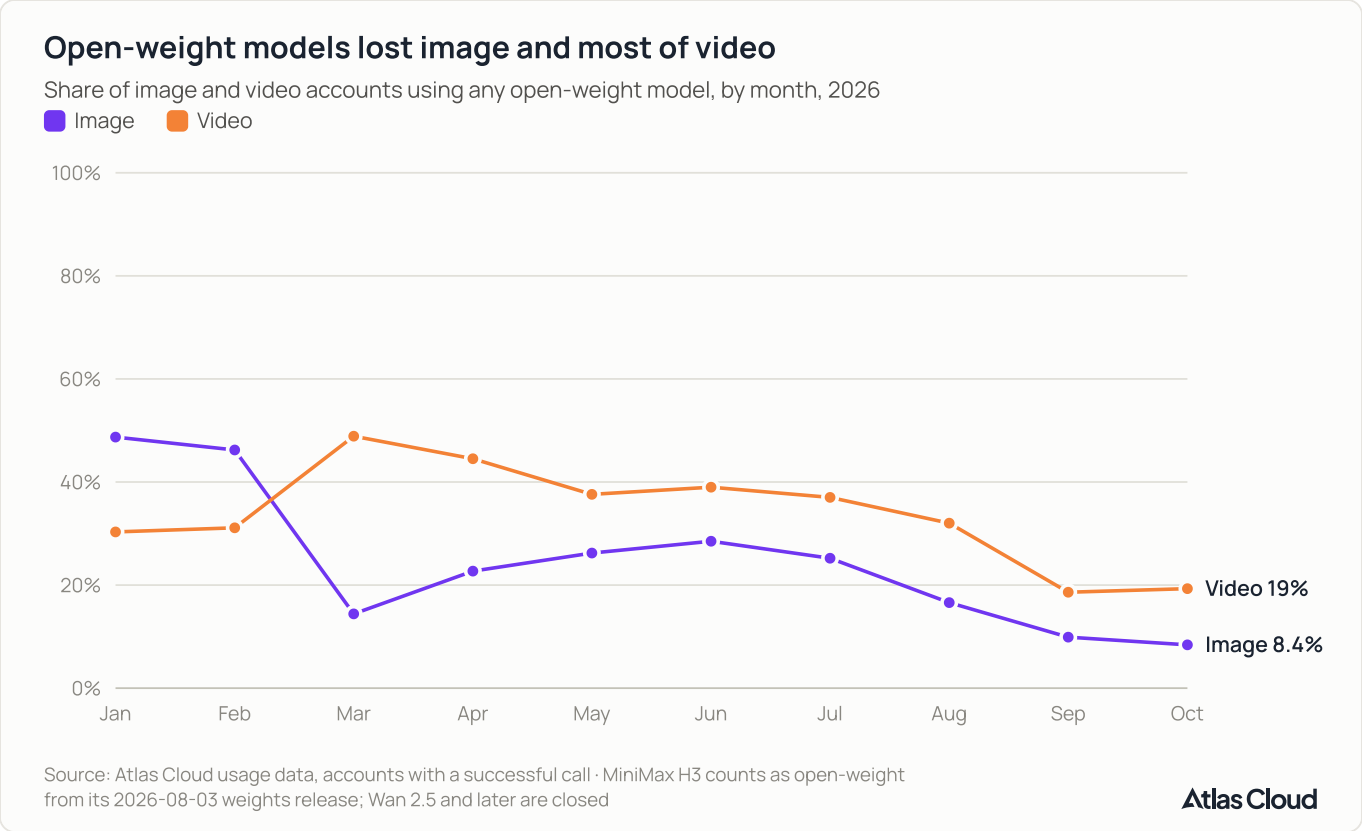
- The image lead changed hands **three times** in 2026, and video once, from Wan to Dreamina Seedance in July.
- In text-to-speech, ElevenLabs passed Grok Voice, which was on 70% of text-to-speech accounts in July.
- In 3D, Tripo went from 5% to **49%** of 3D accounts in three months.
- Builders add the new leader more than they replace the old one: from July to September, 20 to 28% of continuing image and video accounts changed their main family each month, and across the year about **60%** of switchers still used the old family the next month.
- Apart from Wan endpoints Atlas retired in September, every displaced leader stayed available.

## Open weights lost image and video

- Image accounts using an open-weight model fell from 49% in January to **8%** in October, as FLUX gave way to closed models: builders followed the top-rated editing models, which are closed (*our reading*).
- Video fell from a March peak of 49% to **19%**, as Wan's newer versions are closed.

- Closed models also passed open ones in LLMs (see Language models).

FIGURE O.1




## How consumption changed

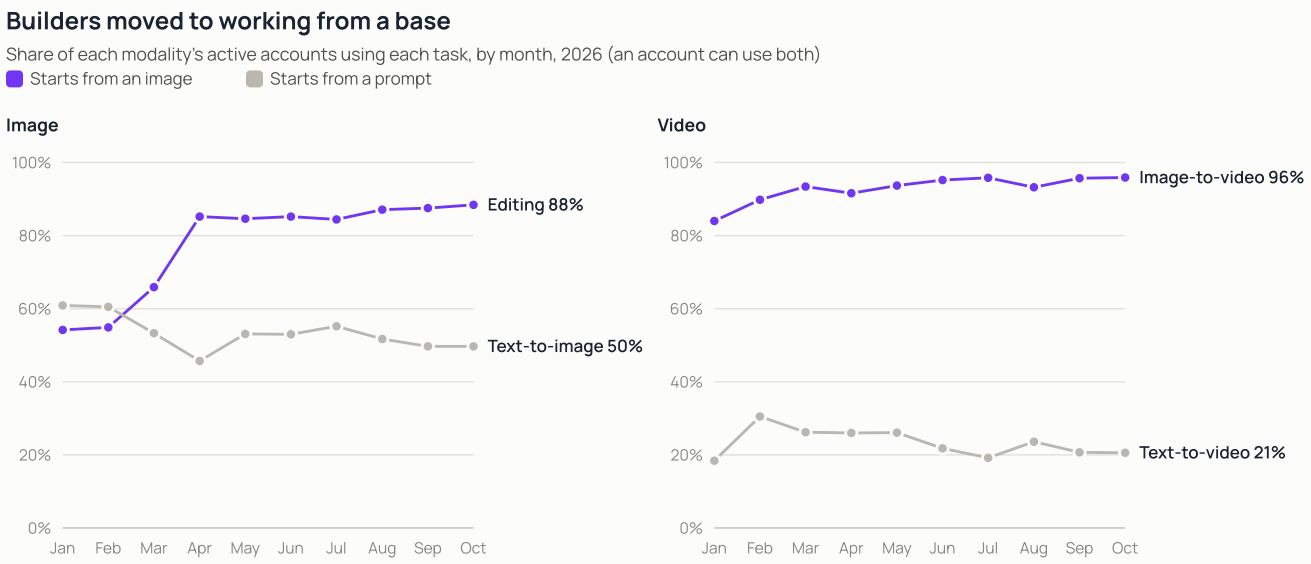
**Builders start from an image and pick models a notch below the top.**

**What it means (*our reading*):** competition is moving to editing and image-to-video quality at the resolutions builders use, and a model wins on its output in the builder's own workflow more than on its rating or price.

### Builders moved to working from a base

- Image accounts that edit rose from 54% in January to 85% in April and **88%** in October, while text-to-image fell from 61% to 50%. Half of image builders still generate from text, and most of those also edit.
- In video, image-to-video went from 84% to **96%** of accounts, and text-to-video ended near where it started, at about a fifth.
- By volume, editing was 80 to 87% of image requests every month (**82%** in October), and **92%** of October's video requests animated an image, so the heaviest users already worked from a base in January and most other builders followed during 2026.

FIGURE B.1

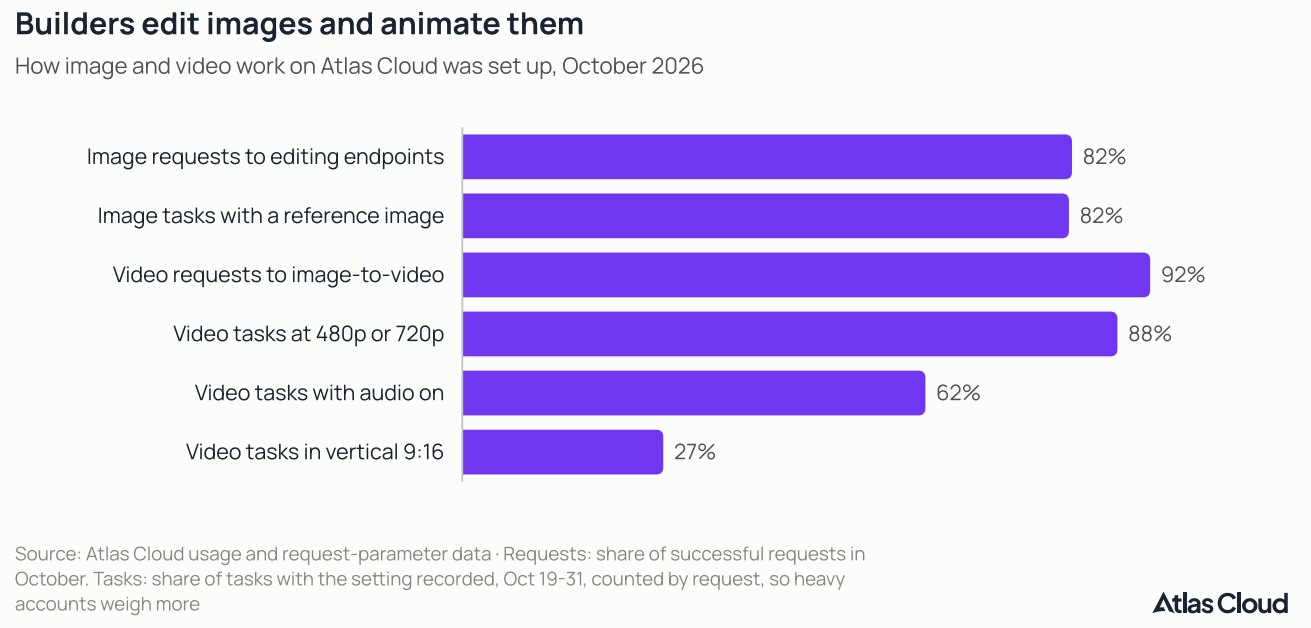



## How the work is set up

A two-week snapshot from October 19 to 31:

- **82%** of image tasks carried at least one reference image, and about four in ten carried two or more.
- In video, 53% of tasks animated a single image, and about four in ten passed two or more, such as a first and last frame (*our reading of what the inputs are*).

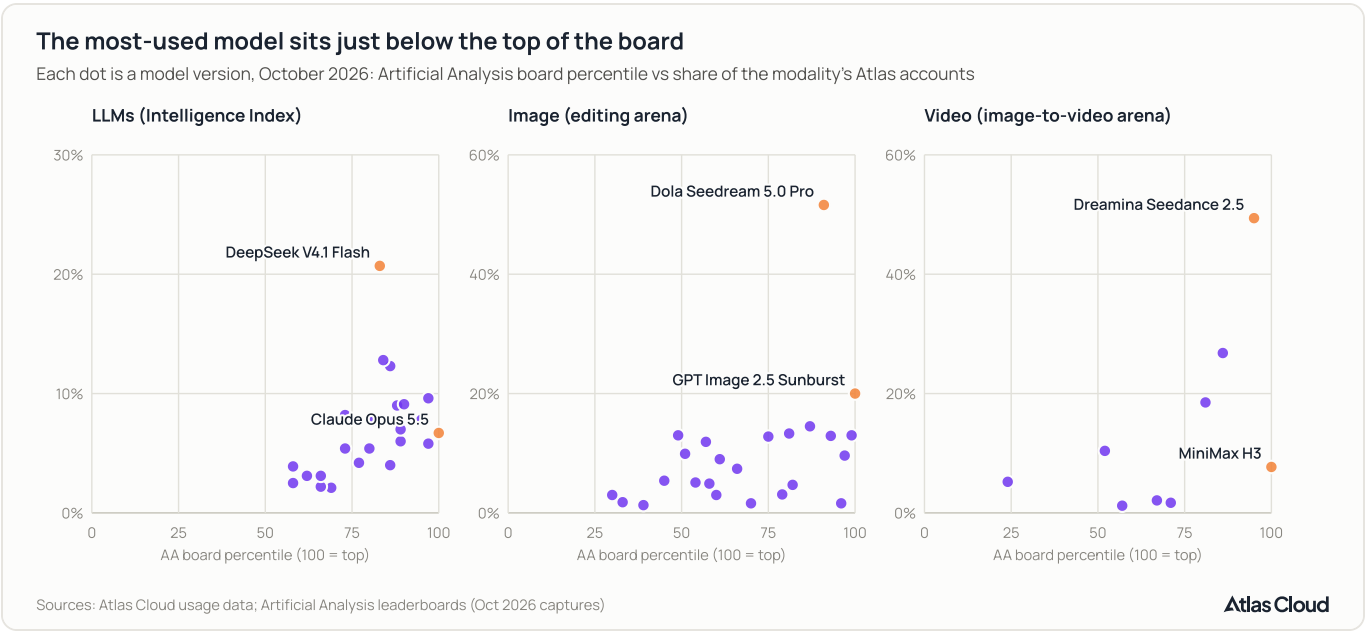
FIGURE U.1



### Do the top-rated models get the most use? No: the most-used model usually sits just below the top

- On **none of the six** Artificial Analysis leaderboards we matched (text-to-image, image editing, text-to-video, image-to-video, text-to-speech and LLMs) was the best-rated model we could match the most used on Atlas in October.
- The most-used model usually sits **just below the top**, between the 80th and 95th percentile of its board:
- **Image editing:** Dola Seedream 5.0 Pro (91st percentile) is on 52% of image accounts; the top-rated GPT Image 2.5 Sunburst is on 20%.
- **Image-to-video:** Dreamina Seedance 2.5 (95th) is on 49% of video accounts; the top-rated MiniMax H3 version is on 8%.
- **LLMs:** DeepSeek V4.1 Flash (83rd) is on 21% of LLM accounts; the top-rated Claude Opus 5.5 is on 7%.
- Ratings still decide who is in the running: models in the top quarter of a board had a median share **two to three times** that of the rest in image and LLMs, under twice in speech, and about ten times in video, where few versions are rated.
- Ratings also count for more than they did on the image and LLM boards: there the rank correlation between rating and use rose from near zero earlier in the year to 0.46 to 0.67 in October. In October it was positive on all six boards, from 0.31 in speech to 0.79 in text-to-video.
    - *Our reading:* builders shortlist from the top of the board, test on their own work, and often keep the near-top model already in their workflow.

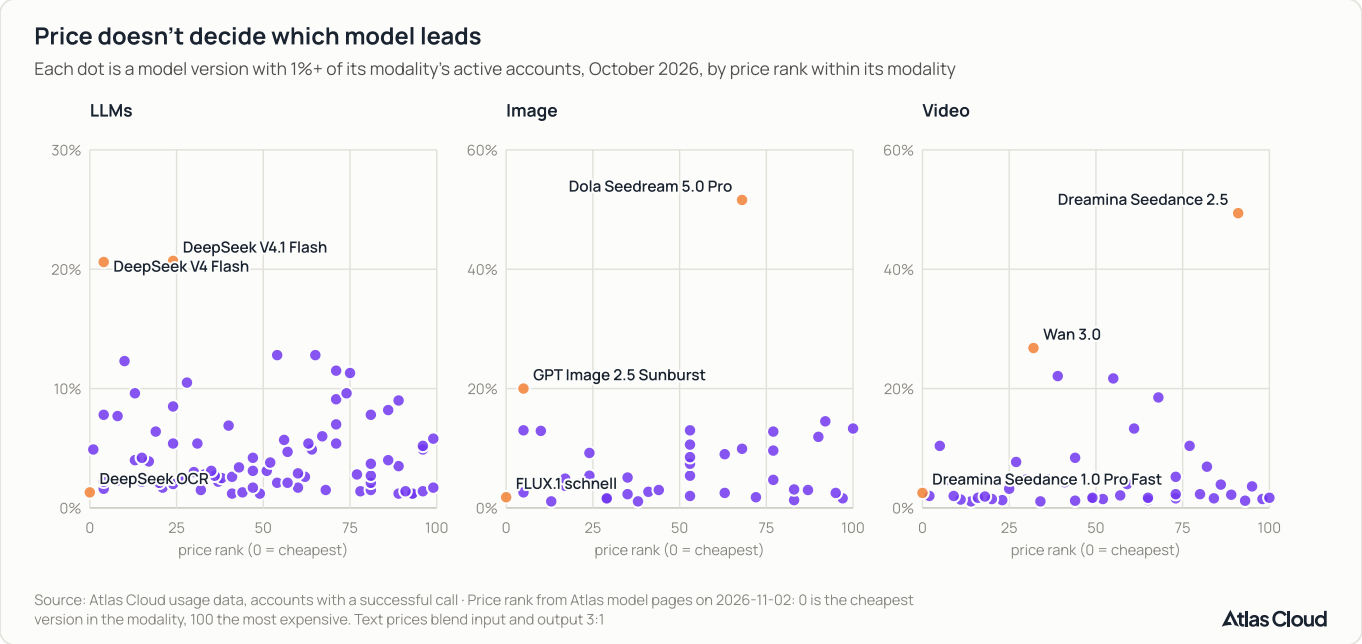
FIGURE 3.2



## How much did price move usage? Little, except for one heavy cut

- In October the rank correlation between a model's price and its share of accounts was 0.06 in video, 0.11 in image and -0.07 in LLMs: **close to no relationship**.
- **The leaders aren't the cheap options.** The most-used video model is priced at the 91st percentile of video models on Atlas and the most-used image model at the 68th. The cheapest image model in wide use, FLUX schnell, fell from 37% of image accounts in January to 2% in October.
- **LLMs are the partial exception:** among models with similar ratings, cheaper ones drew more accounts, and the most-used LLM is priced at the 24th percentile.
- **Price coincided with a usage change once in our data:** a heavy cut with a campaign on two budget video tiers was followed by their daily activity nearly doubling within four weeks, relative to other video families. Routine light cuts didn't reverse declines (see Video).
  - *Our reading:* builders choose media models on output first and price second; price counts for more in LLMs, which they use as a light helper.

FIGURE 2.1



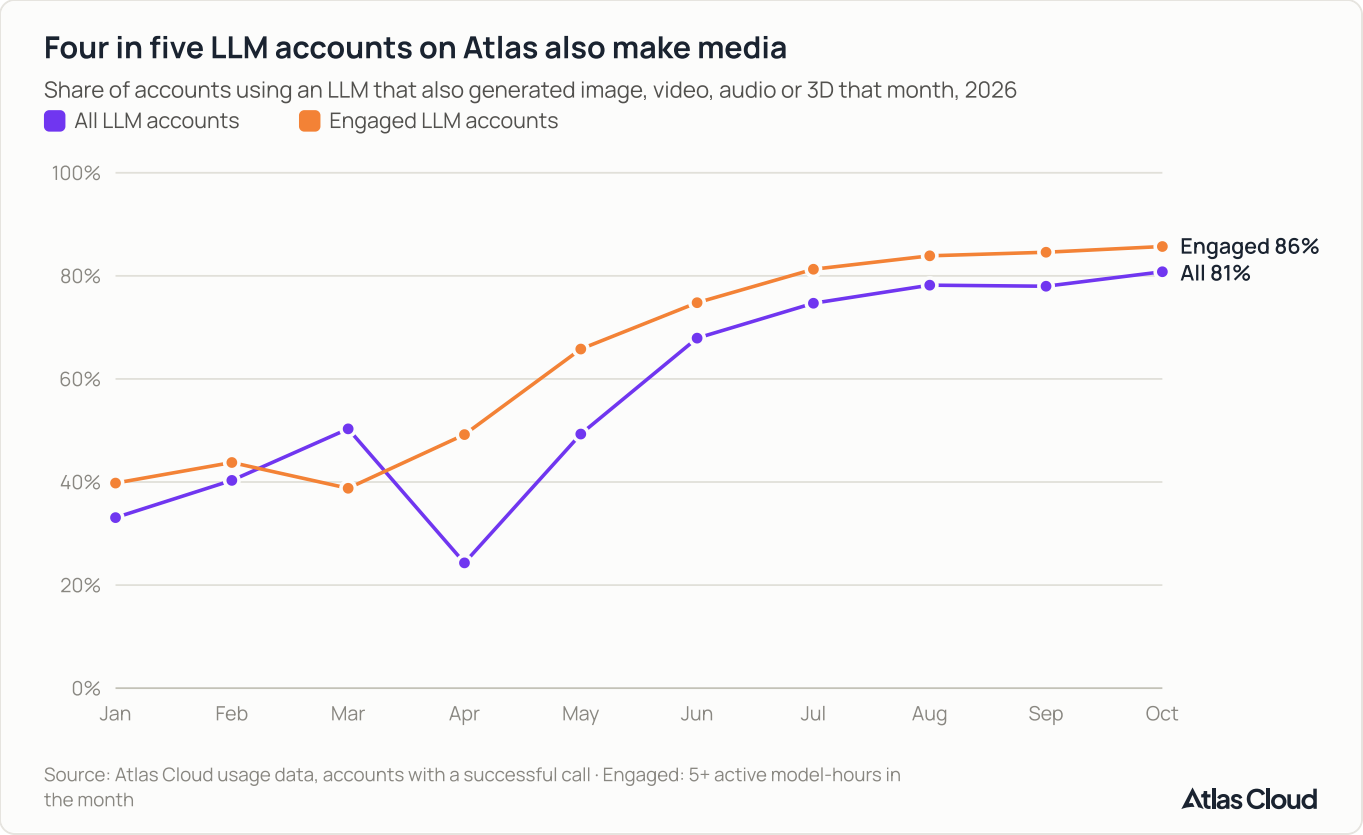

## Most active builders work across media, and come back more often

- **64%** of accounts in steady use worked in two or more modalities in October, up from 35% in January.
- In September, **80%** of accounts using two or more media modalities were active the next month, against 48% of single-modality accounts. *Part of that gap is that they use Atlas more.*

## LLMs get added to media work

**81%** of the accounts that used a language model in October also made media that month, up from 33% in January, and seven in ten of them started on Atlas with media.

FIGURE 1.1




## How business accounts experiment

**Companies moved from testing to steady use; they try models broadly and keep them selectively.**

**What it means (*our reading*):** a model maker wins a company by being kept after the first month.

We use accounts whose owner signed up with a company email as a proxy for companies. *People often sign up with a personal email at work, and a company can have several accounts. Each account counts once however much it uses, so these shares count accounts, not companies, spending or workload.*

### From light use to steady use

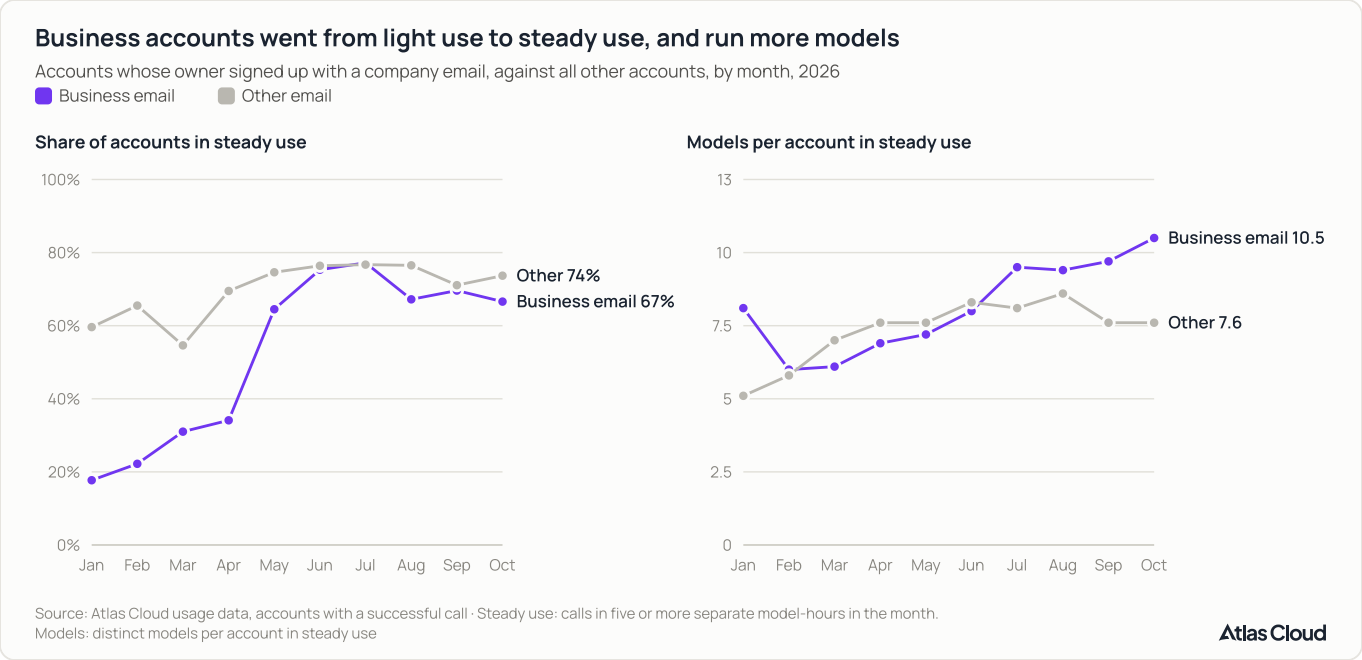
- In January, **18%** of business accounts were in steady use, against 60% of other accounts: most companies were testing with a few calls.
- The number of business accounts in steady use grew about five and a half times from January to October, fastest from April to July. By June and July both groups were at 75 to 77% in steady use, and in October **67%** against 74%.

### Companies in steady use run more models

- Business accounts in steady use ran **10.5** model endpoints in October, against 7.6 for other accounts in steady use, which fits testing several models side by side before choosing (*our reading*).

- Team accounts, set up for several people to share, ran 13 on average. They grew from 0.2% to 1.7% of active accounts, and 74% had two or more members.

FIGURE E.1




## Most first tries aren't repeated, but trying widely goes with staying

- Of the models an account used for the first time in September, 28% were used again the next month by business accounts and 23% by others.
- Of accounts new to Atlas from November 2025 to July 2026, **44%** of those that tried five or more models in their first week were still active three months later, against **8%** of those that tried one.
- *Busier accounts both try more and stay longer, so this is an association, not proof.*

## Business accounts keep growing, and individual builders grew even faster

- Business accounts grew overall from June and ended October about **half above** their January level. Their share of active accounts has risen again since July, from 9% to **13%**.
- Personal-email accounts, which we read as mostly individual creators, prosumers and developers (some are people at work), **more than doubled**, most of it between May and August. That surge took the business share from 19% in January to its July low.
  - *Our reading:* the surge coincides with maturing tools, workflows and skills, as more people experiment and learn to fit AI into their work, and with lower prices that put capable models within reach of individuals. In the heavy-cut case, 67% of Dreamina Seedance 2.0 Mini's new users after the cut were new to Atlas.

# About this report

## What Atlas Cloud sees

Atlas Cloud is an API gateway that sells models as a service: it hosts some models itself and passes calls through to the model's provider for others. Because one account reaches image, video, audio, 3D and language models, we can see how the same accounts combine them. It is a selected slice, mostly media builders and buyers of open-weight models on price; companies committed to OpenAI, Anthropic or Google mostly call those providers directly. We compare with fal, a media-first inference platform, and OpenRouter, a language-model router, where they measure the same thing.

## How we counted

- The unit is an account, not a customer or a person. An account is active in a month when it made at least one successful call to a generative model; Atlas's internal accounts, apps and utilities are left out.
- Every account counts once, whatever it spends. *One company running production work can use more than thousands of individual accounts combined, so account shares can understate how much of the work businesses do.*
- Accounts in steady use made calls in five or more separate model-hours in the month, and the main headlines were rerun on them.
- We give shares and correlations, never account counts, and no account-based figure rests on fewer than five accounts.
- The window is January to October 2026, and every figure is "on Atlas Cloud".
- Request-setting snapshots, the daily price series and the moderation-rejection figure come from model-level aggregates that can't exclude Atlas's internal accounts; every other figure does.
- Request-settings data (shapes, sound and references) covers October 19 to 31 only.

## How we handled the data

- **We used usage records, not content.** Every Atlas usage figure comes from billing and task records (which model an account called, when, for what kind of task, with how many input files and which output settings) and two account attributes used to group accounts: whether the sign-up email is a company address, and whether it is a team account. Benchmark ratings and other platforms' figures come from the published sources we cite.
- **We opened no customer content for this report.** No prompts, uploaded images, generated images, video, audio or text were opened, sampled or used.

- **We publish only aggregates and name no customer in the findings.** Figures are aggregated across accounts and published as shares, with no account counts, and no account-based figure rests on fewer than five accounts. Request-setting figures are shares of successful tasks from October 19 to 31; that extract carries no account field, so we could not run the five-account check on them. No customer is named in the findings.

## Who we are

**The AI Creation Engine for creators and developers.**

Atlas Cloud helps teams create content and build AI-powered products with leading AI models. We bring video, image, audio, and language models together through one API, alongside no-code creation tools and enterprise support. With competitive pricing and enterprise-grade reliability, we serve studios, agencies, brands, and AI product teams.

# Video

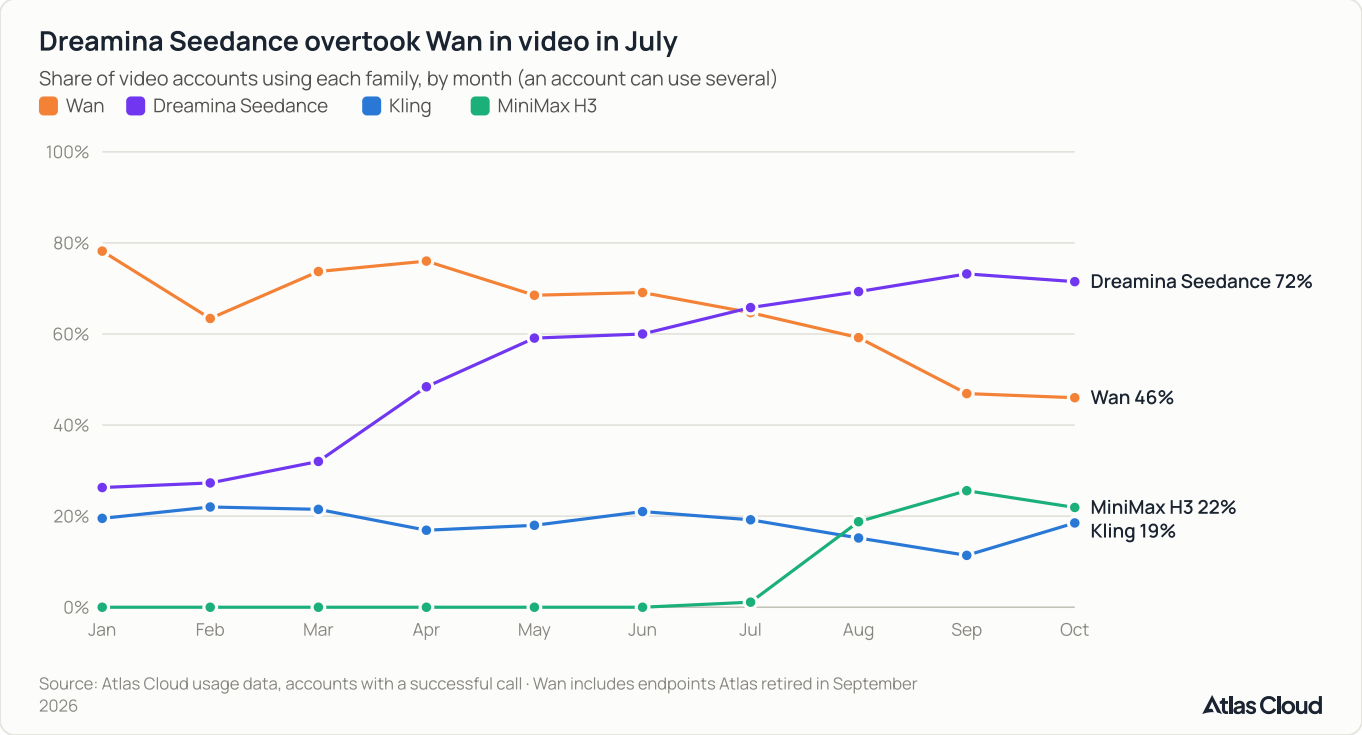
**Builders animate stills; Dreamina Seedance leads by accounts and Wan by volume.**

**What it means (*our reading*):** video models compete on image-to-video quality, and two leaders can coexist by serving different builders.

## Dreamina Seedance leads, Wan produces the most

- In October, Dreamina Seedance was on **72%** of image-to-video accounts and 67% of text-to-video accounts, with Wan second (45% and 34%) and MiniMax H3, launched at the end of July, third at about a fifth.
- By seconds of video generated, **Wan led every month of 2026**: fewer accounts, more volume each. Wan's figures include endpoints Atlas retired in September.

FIGURE V.1




## A late-October snapshot: short clips with sound

Our request settings cover only October 19 to 31, so *read these as a two-week snapshot, not a trend.*

- 480p and 720p made up **88%** of video tasks and 1080p under 4%, and audio was on in **62%**.
- Vertical 9:16 edged out 16:9 (27% against 24%), and the most common lengths were automatic, 15 seconds and 5 seconds.

- These are the settings of clips for social feeds and ads, or of cheap drafts before a final render (*our reading*).

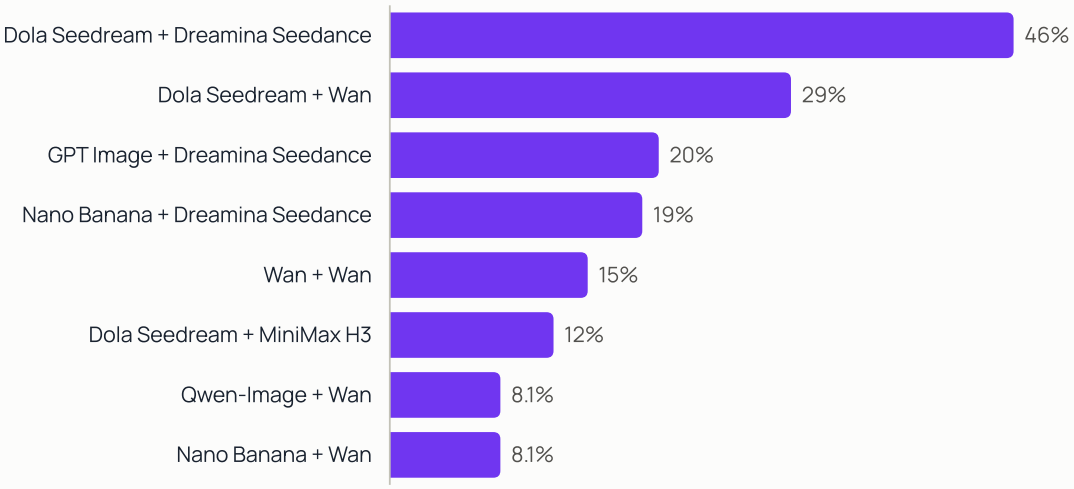
## Image and video come from the same builders

- **62%** of image-to-video accounts also made images in October.
- Half used an image model and an image-to-video model in the same hour at least once, most often Dola Seedream with Dreamina Seedance. That fits making a still and animating it in one sitting, *but hourly data can't show whether one call feeds the next.*

FIGURE U.2

**Dola Seedream with Dreamina Seedance is the most common image-and-video pair**

Pairs of image and image-to-video families used by the same account in the same hour, October 2026




Source: Atlas Cloud usage data, accounts with a successful call · Base: accounts that used an image model and an image-to-video model in the same hour at least once (50% of image-to-video accounts); an account can show several pairs. Same-hour use isn't proof of a pipeline. Atlas resells and promotes both Dreamina models

Atlas Cloud

## Launches that last bring users back

- **49%** of the first-week users of Dreamina Seedance 2.5 came back in week two, and **45%** of the accounts that tried it in its first month were still using it the month after.
- For MiniMax H3, Kling 3.0 Pro and Dreamina Seedance 2.0 Fast, those figures were 15 to 21% and 16 to 29%.
- Dreamina Seedance 2.5, at the 95th percentile of Artificial Analysis's image-to-video board, is **the only one of them that leads**.

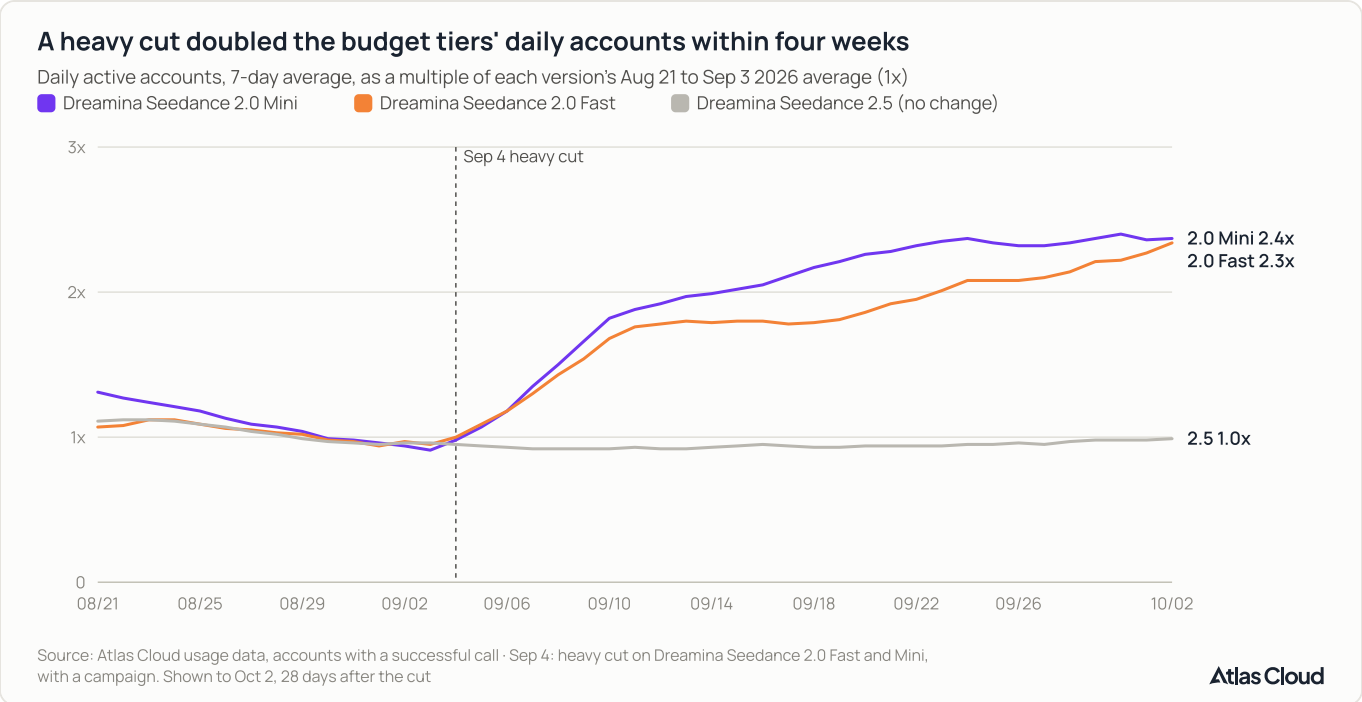
## The heavy-cut case

**What it means (*our reading*):** a heavy cut with a campaign can draw new buyers to a budget tier while the flagship holds level; routine light cuts came after declines and didn't reverse them.

On 2026-09-04, Atlas gave Dreamina Seedance 2.0 Fast and Dreamina Seedance 2.0 Mini a heavy cut, with a campaign.

- Over the next 28 days, relative to other video families (the comparison group excludes the cut family), their daily active accounts (counted per endpoint, so an account using both tiers counts twice) rose **96% and 99%** and their requests about three and four and a half times.
- Dreamina Seedance 2.5, whose price didn't change, held level.
- Figure 2.3 shows the daily series before that adjustment.

FIGURE 2.3



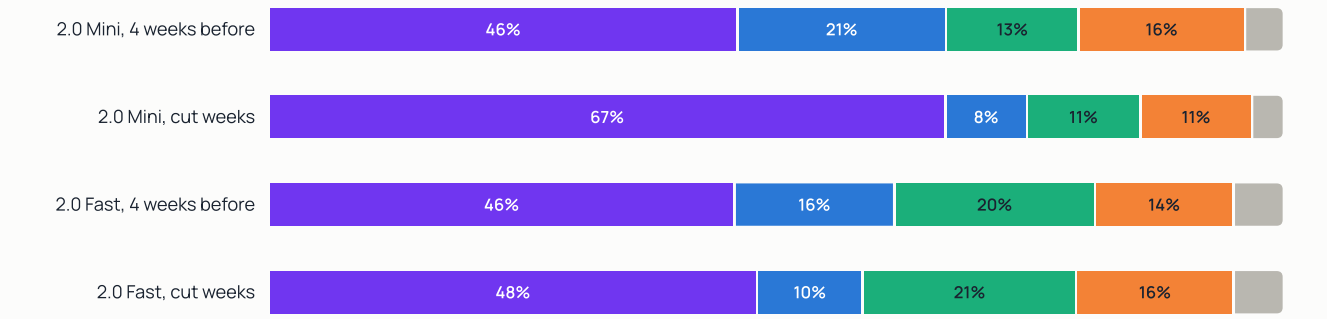

### Who came, and how routine cuts compare

- Mini drew 2.2 times as many first-time users in the four weeks after the cut as in the four weeks before, and **67%** of them were new to Atlas, against 46% before (Figure 2.6). Fast drew about the same mix as before.
- Routine price changes, mostly light cuts, followed declines: accounts fell a median 17% relative to other families in the two weeks before and were still **13% lower** in the two weeks after (Figure 2.2).

FIGURE 2.6

## The cut drew accounts new to Atlas to Mini; Fast drew the same mix as before

Where first-time users of Dreamina Seedance 2.0 Mini and Fast came from: the 4 weeks from the Sep 4 2026 cut vs the 4 weeks before



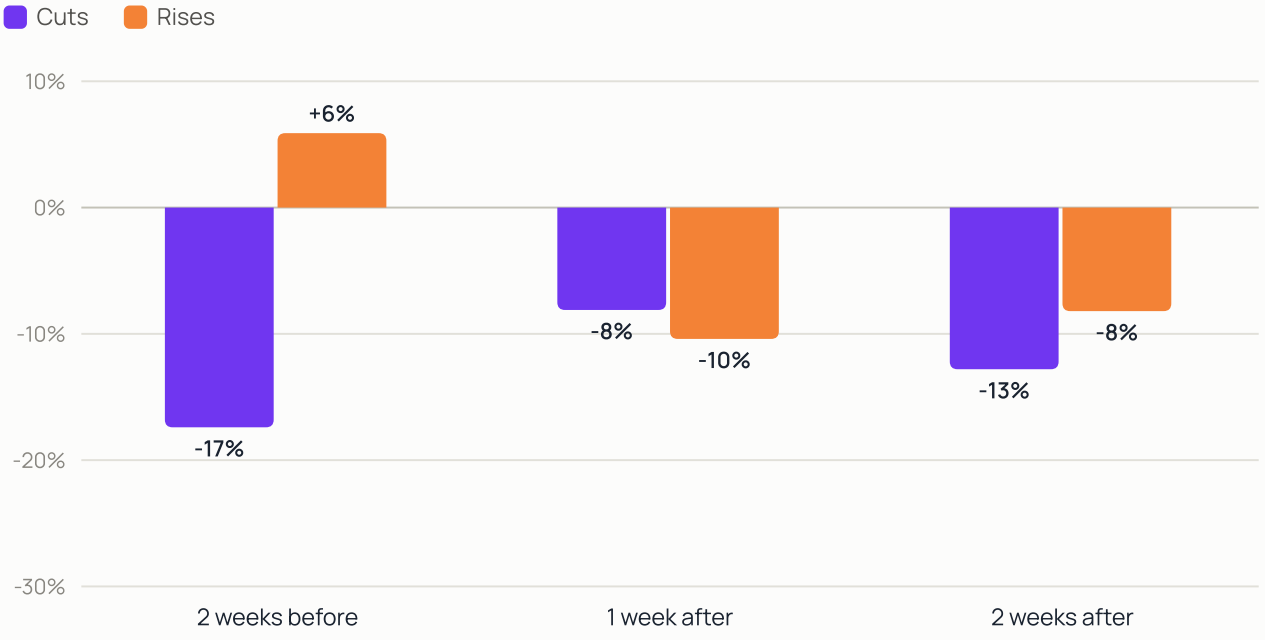

Source: Atlas Cloud usage data, accounts with a successful call · Mini drew 2.2 times as many first-time users in the cut weeks as in the weeks before, Fast 1.6 times. 'From' is the account's main video family the month before

Atlas Cloud

FIGURE 2.2

## Routine list-price cuts typically came after declines

Median change in a model's daily accounts relative to other families, around list-price changes




Source: Atlas Cloud usage data, accounts with a successful call · Routine changes to one model, mostly light cuts, 28+ days after launch, no other change within 14 days

Atlas Cloud

### How sure we are

One heavy cut is a case study, and the result measures price and campaign together. The routine-change test uses list prices, which miss short discounts, and its samples are small. Atlas set these prices and serves these models.

## How it compares

- fal's August 2026 shelf lists three of Atlas's four most-used video versions (fal Volume 2).
- On OpenRouter, the week of October 26, video requests went 26.1% to Dreamina Seedance, 19.7% to HeyGen and 18.5% to Google (OpenRouter rankings).
- Neither has published a price experiment that we found.

# Image

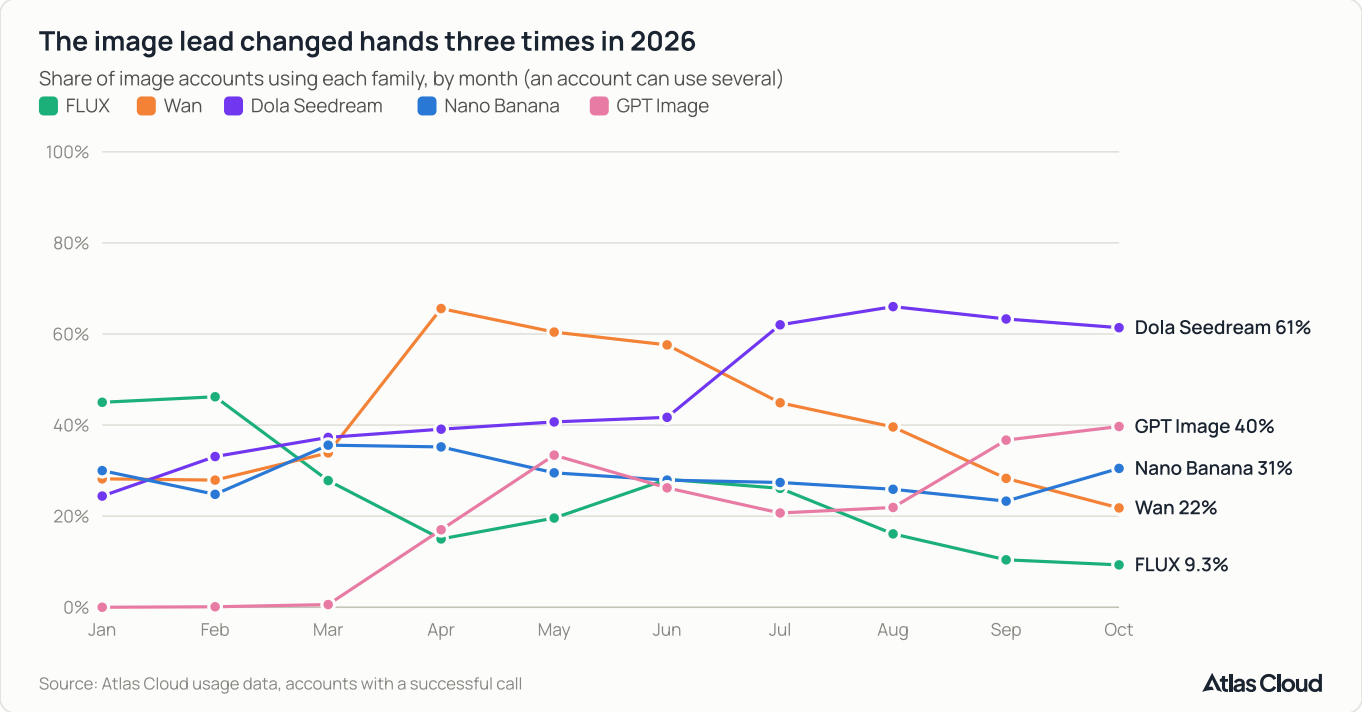
**Image builders edit with references, and the leader faces a fast-rising challenger.**

**What it means (*our reading*):** the image contest is about editing with reference images, and a top-rated launch can take share within a quarter.

## Dola Seedream leads, GPT Image is closing in, and companies lean to GPT Image

- In October, Dola Seedream was on **63%** of image-editing accounts and 52% of text-to-image accounts.
- GPT Image **nearly doubled** its share of both since July, to 36% and 46%, with Nano Banana third.
- Among accounts using image models, GPT Image was the main family for **41%** of business accounts against 22% of others, and Dola Seedream for 27% against 48%.
- FLUX led until February, Dola Seedream briefly in March, Wan from April to June and Dola Seedream again from July. Each displaced leader was still used by 28 to 45% of image accounts in the month it lost the lead.
- 52% of Dola Seedream 5.0 Pro's first-week users came back in week two, the highest of the image launches we tracked.

FIGURE I.1



## Editing, with references, at 2K

- **2K** was the most requested output size from October 19 to 31, and the 4K setting was under 7%: enough for product pages and social posts without paying for 4K (*our reading*).
- In the same two weeks, portrait shapes (taller than wide) were requested about **1.5 times** as often as landscape: 26% of image tasks against 17%, with half setting no shape. That fits images made for vertical feeds (*our reading*).
- The average image account used about four image models in October, and **47%** used two or more image families, which fits keeping a second model to compare or fall back on.

## How it compares

- fal's most-requested image models in late 2025 were Nano Banana Pro Edit, Reve, Dola Seedream 4.5, Qwen Image Max and FLUX.1 dev (fal Volume 1); Atlas at the same point was mostly FLUX.
- On OpenRouter, image requests in the week of October 26 went 47.5% to Google and 28.3% to OpenAI (OpenRouter rankings), likely skewed to LLM developers adding images (*our inference*).

# Audio

**Builders add audio to image and video work, try it fast and keep it rarely.**

**What it means (*our reading*):** audio models compete for a slot inside image and video work, where video models with native audio are also rivals.

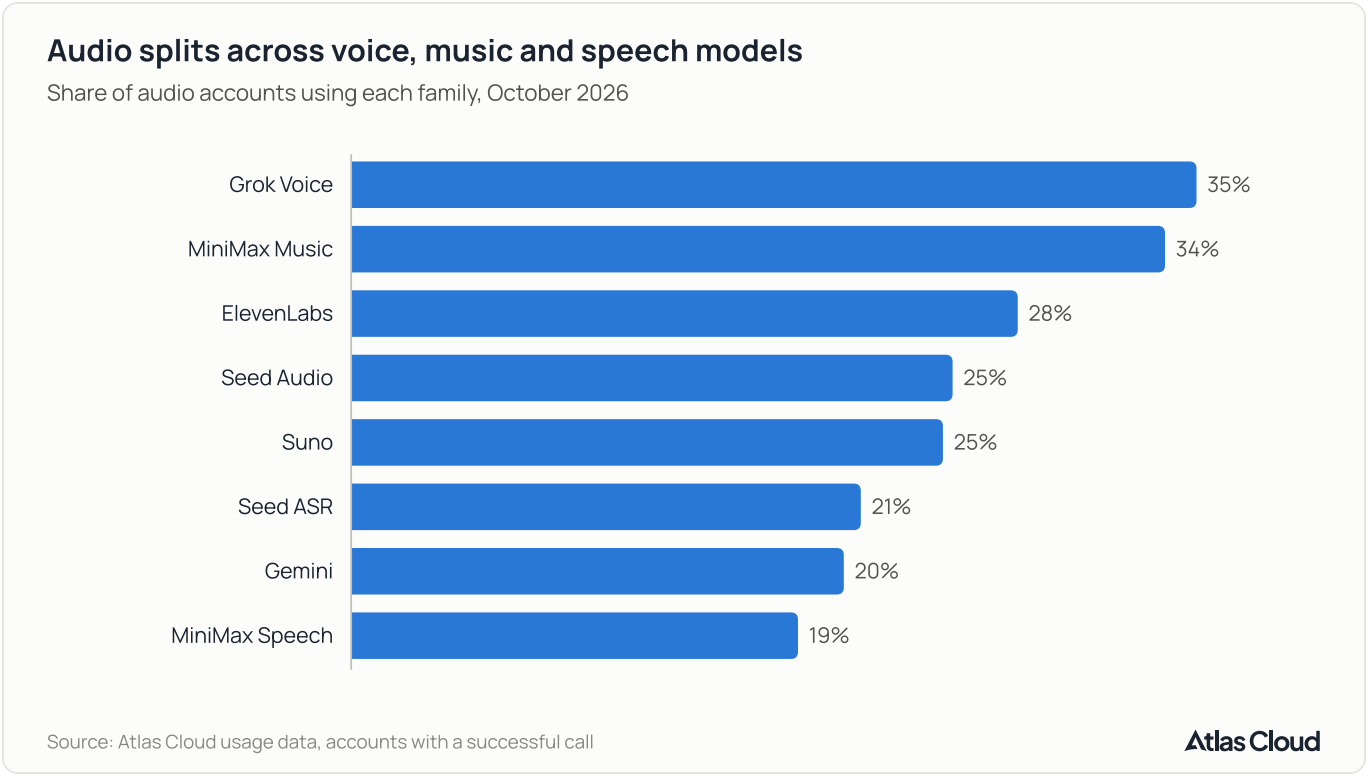
## Audio is an add-on

**95%** of audio accounts also used another modality that month, the highest of any modality, and half used two or more audio families. 94% used a closed model.

## No single leader across audio

- In text-to-speech, ElevenLabs (**47%** of text-to-speech accounts) passed Grok Voice (40%), which had 70% in July, and Gemini more than doubled its share.
- In music, MiniMax Music (68% of music accounts) and Suno (50%) held steady, and many music builders use both.
- Audio buyers are still shopping: from July to September about **four in ten** continuing audio accounts changed their main family each month, roughly twice the rate in image and video.

FIGURE A.1



## Tried fast, kept rarely

- Seed Audio 1.0 and ElevenLabs v3 reached a quarter of audio accounts within two and six days of launch, on the small audio base of June, but only **17% and 7%** of their first-week users came back in week two.
- **59%** of September's audio accounts made no audio calls in October, against 34% for image.

## How it compares

- fal's most-requested text-to-speech models in late 2025 were from ElevenLabs and MiniMax.
- On OpenRouter, speech requests in the week of October 26 went 54.0% to Google, 15.9% to ElevenLabs and 11.6% to Fish Audio (OpenRouter rankings).
- ElevenLabs ranks high on all three platforms.

---

# 3D

**3D is small on Atlas, but the builders who use it use it more each month.**

**What it means (*our reading*):** 3D is a modality to watch, not yet one to call. The builders who adopt it go deeper rather than more of them arriving, and part of the small base is our thin lineup. A 3D asset can be lit, posed and rendered from any angle, again and again, where a video clip is a fixed series of frames, which is one reason the use may keep deepening. World models, which generate explorable environments rather than single objects, are a related next step.

## A thin lineup, and we say so

- Atlas added 3D in June and lists 14 3D endpoints from five families (Seed3D, Hunyuan3D, Tripo, Meshy and Hi3D), a small share of its 400+ models. All five are closed; we host no open-weight 3D model.
- 3D has held at about **3%** of active accounts since June. *With this few models, part of that is our supply rather than demand.*

## Builders went deeper

- The average 3D account made about **4.7 times** as many successful 3D requests in October as in June, while the number of 3D accounts grew only modestly.
- **92%** of 3D accounts also used another modality in October. Every family we host takes an image as input, so 3D comes after image work: make the reference, then build the model (*our reading*).

## The lead moved fast

- Seed3D 2.0 led from June to August and was on 69% of 3D accounts in June. Tripo H3.1 led from September and was on **49%** in October, and Meshy 7 reached 36%.
- 38% of 3D accounts used two or more 3D families in October, up from 15% in June: builders compare models here too.

## The leading 3D model costs more per output than the median image model

The leading model, Tripo H3.1, lists at $0.10 per 3D model on Atlas, about **three times** the median image price, and Seed3D 2.0 at $0.35.

# Language models

**On Atlas, LLMs are a light helper for media builders, and closed models are gaining.**

**What it means (*our reading*):** media builders buy LLMs bundled with media work and pick low-cost models over frontier ones. *Read this chapter as a slice: buyers committed to OpenAI or Anthropic mostly go direct, so closed frontier models are under-represented here.*

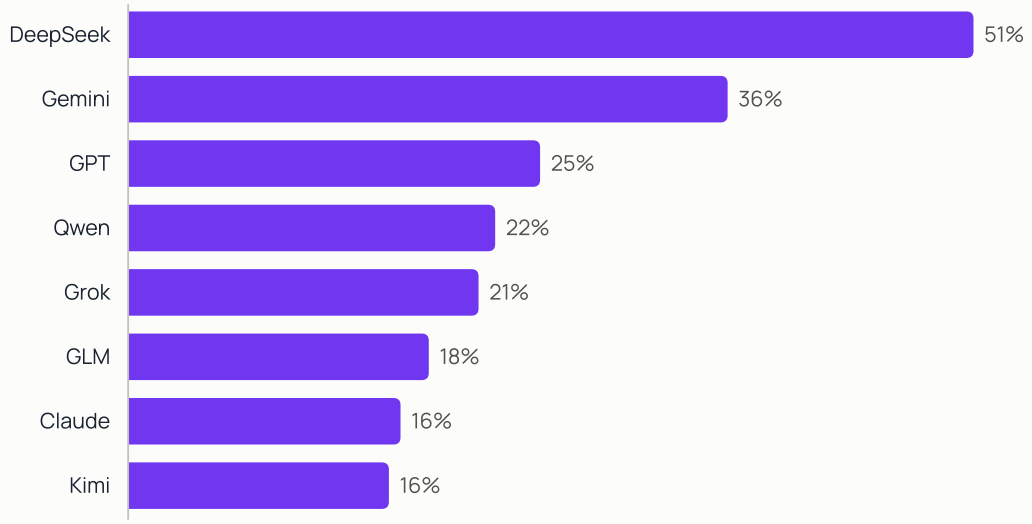
## DeepSeek leads, and versions churn

- DeepSeek was used by **51%** of LLM accounts in October, ahead of Gemini at 36%, GPT at 25%, Qwen at 22% and Grok at 21%, and led every month of 2026 except April.
- The leading version changed **five times** in 2026, to DeepSeek V4.1 Flash in October, at the 83rd percentile of Artificial Analysis's Intelligence Index. Claude Opus 5.5, at the top, is on 7% of LLM accounts.

FIGURE L.2

### DeepSeek leads Atlas's slice of the LLM market

Share of LLM accounts using each model family, October 2026



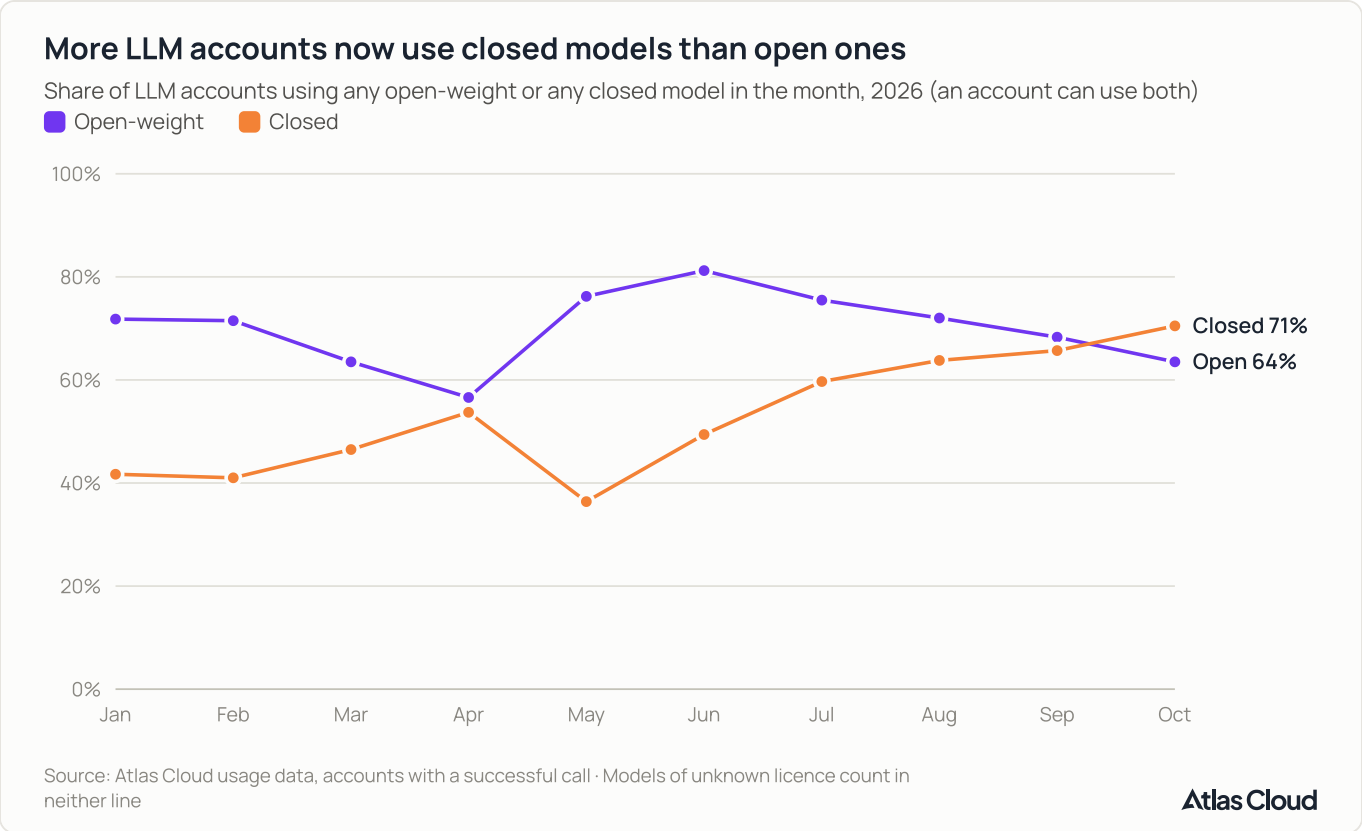

Source: Atlas Cloud usage data, accounts with a successful call · Atlas sees LLM traffic that comes through a multi-model API; frontier-model buyers mostly go direct

Atlas Cloud

## Closed models passed open ones

In October, **71%** of LLM accounts used a closed model and 64% an open one, against 42% and 72% in January, as media builders who joined during the year reached for Gemini, GPT and Grok as well as DeepSeek. Models of unknown licence count in neither line.

FIGURE L.1




## Where the LLM-and-media rise came from: mostly new, media-first accounts

- **62%** of October's LLM accounts first used Atlas between July and October, and **85%** of those also made media.
- Of October's accounts using both, **71%** started on Atlas with media, 24% with both on the same day and 5% with an LLM.

## LLM use is light

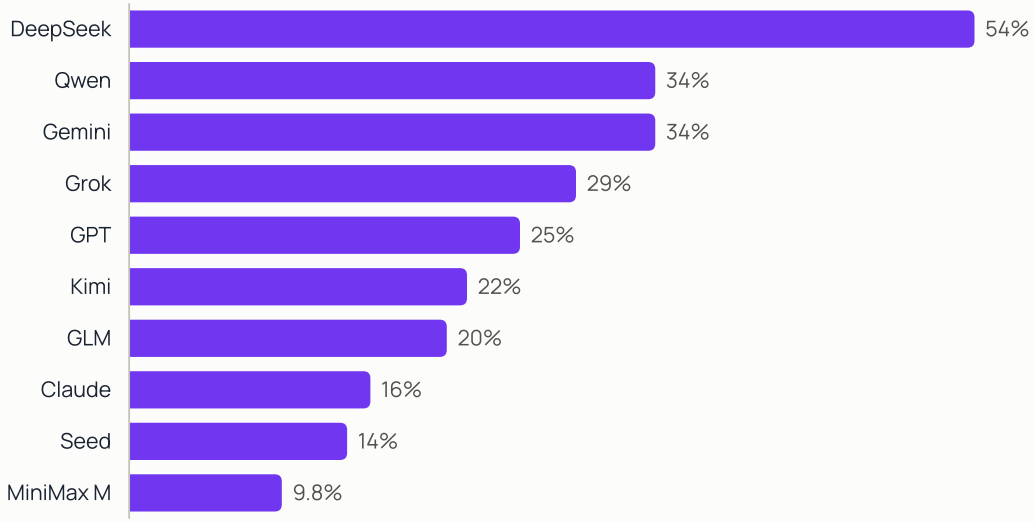
- The median account using both made **28 text calls** in October, and from June to September fewer than half used both again the next month.

- They likely use an LLM for scripts, captions or prompts (*our inference; we opened no prompts for this analysis*). From July to October they reached for DeepSeek (54% of them), Gemini and Qwen (34% each), then Grok and GPT (Figure 1.4).
- The 81% holds on accounts in steady use (86%). Text volume is concentrated, so this chapter counts accounts, never tokens.

FIGURE 1.4

**DeepSeek leads the LLMs media accounts reach for, then Gemini and Qwen**

Share of accounts using both an LLM and media that used each LLM family, Jul to Oct 2026




Source: Atlas Cloud usage data, accounts with a successful call

Atlas Cloud

## How it compares

- OpenRouter reported in September 2026 that its open-weight share of tokens passed 50% for good in April, peaked at 79% in early July and stayed above 62% after that (OpenRouter). It counts tokens, which heavy open-model users dominate, and we count accounts.
- fal writes that generative media, LLMs and agents have "collapsed into one loop"; on Atlas the same accounts use both, though our hourly data can't show a loop.

# What it means by industry

Usage data has no industry field. This chapter pairs the patterns in this report with the business customers our team works with directly, ordered from the largest group to the smallest. *Read it as our interpretation; each point names the finding it rests on.*

## App builders and software platforms

- **The largest group of business customers:** creative and editing apps, AI companions and games, and content tools. Most resell generation to their own users, so they feel model and price changes first.
- Leaders change and builders keep the old model, so apps need routing and fallback across models, not one integration per model maker.
- 81% of LLM accounts also make media, and seven in ten of them started with media: the text stack gets added to the media stack, so both end up in the same product.
- LLMs are the one place where lower prices go with more share, which favors low-cost models for high-volume helper and agent tasks.

## Film, TV, streaming and short drama

- **The second-largest group:** film studios, short-video platforms, short-drama producers and digital publishers adapting stories to video.
- Production-heavy video use looks like Wan's: fewer accounts, each generating far more seconds.
- Video tasks that pass two or more inputs, such as a first and last frame, point to the control studios need over shots and continuity.

## Brands, retail and e-commerce

- **Mostly apparel and footwear makers and online marketplaces:** companies with large catalogs and new products every season.
- Product images in, versions out. Editing (88% of image accounts), edits that combine several reference images, and image-to-video fit a catalog-scale workflow: place the product in new scenes, then animate it for each placement.
- The library of approved product and brand images becomes the asset that matters most, because most versions start from one.
- **Where agencies fit.** *Our reading:* with brands running the same models in-house, the agency role shifts to the idea, the judgment of what is good enough, and building the workflows that produce the versions.

## Enterprise and vertical software

- **Companies in insurance, life sciences, IT operations, energy and telecom.** Most start on language models, so the LLM findings apply most: lower prices go with more share, and the leader changes within a quarter.
- Media comes second for them. The same move from text to text plus media that the 81% shows is where we expect their next spend, in documents, training material and customer communication (*our reading*).

## Creators and small studios

- **Many arrive as individuals rather than companies.** Personal-email accounts more than doubled over the year, most of it between May and August.
- **Many are experimenting on projects of their own**, from short films to social channels, and some are starting influencer accounts. The late-October settings fit that: vertical 9:16 edged out 16:9, 15- and 5-second clips were among the most common lengths, and most clips had sound (*our reading*).
- Individual builders are likely the most responsive to price (*our reading*): the heavy cut on a budget tier drew mostly users new to Atlas.
- They are a growing supply of AI production skill that brands and studios can hire, wherever the talent sits.

# What it means

## For teams choosing models

- **Test the task you run;** on Atlas that is mostly image editing and image-to-video.
- **Expect the leader to change within a year:** shortlist from the top of the leaderboard and test on your own work.
- **Keep a fallback for moderation rejections.** Recorded rejections were about 5% of image and video requests on Atlas in October; how providers code them varies.

## For model makers

- **Demand is image editing and image-to-video,** with audio bought alongside.
- **Launches that last bring users back in week two,** and the leading launches drew mostly builders new to the task.
- In our one case, **a heavy cut with a campaign on a budget tier coincided with new accounts arriving** while the flagship held level.

## What we'll track next

- World models, once they show up in Atlas usage.
- Whether audio keeps growing inside video stacks.
- Whether agents start chaining image and video calls, which needs finer timing than hourly data.

# Sources


- fal, The State of Generative Media, Volume 1, February 2026, and Volume 2, September 2026.
- OpenRouter, rankings (image, video and speech, week of October 26, 2026), The Token State of Play (September 22, 2026) and DeepSeek V4 Is Earning Agentic Token Share (June 30, 2026).
- Artificial Analysis leaderboards: text-to-image, image editing, text-to-video, image-to-video, text-to-speech and the Intelligence Index, monthly captures, 2026.

# About the author

Colin Mo is Head of Product Marketing at Atlas Cloud. If you want to talk about the findings, connect on LinkedIn.

- **Questions about the report, press and partnerships:** marketing@atlascloud.ai
- **Running these models at scale:** sales@atlascloud.ai

# About Atlas Cloud

Atlas Cloud is the model gateway for teams that build with generative media. One API key reaches 400+ image, video, audio, 3D and language models, including nearly every model in this report (a few endpoints were retired during the year), with OpenAI-compatible endpoints for LLMs.

- **Priced to scale.** Pay-as-you-go prices, set competitively and published on every model page. Lindy moved its AI agents to DeepSeek V4 Flash on Atlas Cloud and cut inference costs on the migrated routes by about 90%.
- **Skip the procurement stage.** One account and one bill cover every model, so adding a model means changing a line of code, not onboarding a new vendor. Yapper adds frontier models through one integration, with no new vendors.
- **Service that stays with you.** Enterprise accounts get a dedicated channel and a named engineer. Versatile Media's MOKE runs its film pipeline on Atlas Cloud, from script to final cut.

Try the models in this report at atlascloud.ai.

# Atlas Cloud

Image, video, audio, 3D and language models through one API key and one bill.

**atlascloud.ai**