State of AI Dubbing 2026 · Mid-Year Update · Version 2.0 · September 2026

State of AI Dubbing 2026, mid-year update: where the paid minutes go

Between January 2025 and August 2026, Perso Dubbing recorded 386,068 platform projects of every type and status. The June edition mapped the first 316,856 of them by volume; this update follows the ones that were paid for: 66,846 projects and 303,241 dubbed minutes between December 2025 and August 2026, across 2,902 paying creators, and asks what they localize, into which languages, and from where.

386,068
platform projects, all types and statuses (Jan 2025 to Aug 2026)

303,241
paid dubbed minutes (Dec 2025 to Aug 2026)

2,902
paying creators

68
target languages · 1,162 active pairs (all platform projects, Jan 2025 to Aug 2026)

119
billing countries with a paid subscription (all subscription records to Sep 7, 2026)

82.5%
of paid minutes from the top 10% of paying creators

Taeksoon Kwon (CTO) and the Perso Dubbing Data Team · · perso.ai/research/state-of-ai-dubbing-2026/ · Originally published June 4, 2026

Current edition · Mid-Year Update · September 2026 · Version 2.0

State of AI Dubbing 2026, mid-year update: where the paid minutes go

Between January 2025 and August 2026, Perso Dubbing recorded 386,068 platform projects of every type and status. The June edition mapped the first 316,856 of them by volume; this update follows the ones that were paid for: 66,846 projects and 303,241 dubbed minutes between December 2025 and August 2026, across 2,902 paying creators, and asks what they localize, into which languages, and from where.

M0 · What changed since June, and why this update reads differently

Two platform changes shape this update, and both change what the data can and cannot say. On December 1, 2025 Perso Dubbing moved from unlimited plans to usage-based pricing, so project counts before and after that date are not the same unit. On June 11, 2026 the platform stopped enabling the share link by default, which ends the June reading of the "share rate" as a behavioral signal. Both changes, with the figures, are set out in Chapter M7.

So this update changes the unit of analysis. Instead of counting projects, it follows paid dubbed minutes: the length of translated video that a paying creator produced. Minutes do not depend on the pricing model, do not depend on a share toggle, and are the unit competitors and buyers already use.

Everything else stays. The industry taxonomy, the professional-creator rule, the confidence-interval convention and the "within Perso Dubbing's data" scope discipline are unchanged. Where a June finding does not survive the paid lens, this update says so in Chapter M6 and shows both numbers. Where a June figure could not be reproduced from the June raw data, the methodology lists it.

M1 · Platform scale, cumulative

Between January 1, 2025 and August 31, 2026, Perso Dubbing recorded 386,068 platform projects of every type and status (including audio-only and cancelled projects) across 71 source and 68 target languages, forming 1,162 active source-to-target pairs. The June edition counted 316,856 projects, 36 source and 34 target languages and 909 pairs; the language expansion of 2026 added 34 target languages, most of them long-tail (Urdu, Bengali, Telugu, Persian, Marathi, Malayalam and others).

Cumulative, Jan 2025 to Aug 2026	Value
Projects	**386,068**
Unique creators	**172,688**
Professional creators (6+ projects)	**4,290**
Single-use creators (exactly 1 project)	**144,983**
Categorized projects (Oct 2025 to Aug 2026)	**173,048**
Source × target languages	**71 × 68**
Active language pairs	**1,162**
Share-link-enabled rate (June: "share rate"), Jan 2025 to Jun 10, 2026	**95.9%**
Billing countries with a paid subscription, all subscription records to Sep 7, 2026 (not limited to the data period)	**119 (68 active)**

These are the cumulative figures of the update. Chapters M2 to M5 are about the paying cohort in the usage-based era; Chapter M6 also uses the cumulative and free cohorts, because re-testing the June findings requires the populations the June edition measured.

Reading the two editions together: the June figures of 4,023 professional creators and 2.43 languages per professional creator were computed under a rule the June pipeline did not document; the same rule set that produces 4,290 and 2.39 here gives 3,743 and 2.39 for the June period (methodology, Section 6.4). The June "80+ countries" has no recoverable definition; 119 is the count of billing countries with a paid subscription lifecycle and is not comparable to it. The count runs over every subscription record up to September 7, 2026, including subscriptions that began before January 2025, so it describes the whole customer base and does not correspond to the data period.

Source: `headline-statistics-2026.csv`, `industry-deep-dive-2026.csv`, `per-language-deep-dive-2026.csv`, `v1-period-recompute-2026.csv`.

M2 · The paying creator

The paid cohort is every project created between December 1, 2025 and August 31, 2026 on a Starter, Creator, PRO, Team or Enterprise plan, after removing multi-account abusers, internal and demo workspaces, and cancelled projects. That is 66,846 projects, 2,902 creators and 2,916 workspaces, and 303,241 dubbed minutes.

Two facts frame the rest of the report.

Paid projects are half the platform's projects and almost nine tenths of its minutes. In the same window, projects with a known plan split 49.8% paid and 50.2% free, but paid projects carry 88.7% of dubbed minutes. Most of that gap is structural: the free plan caps output at about one minute (among free projects with a duration, the 99th percentile of length is 1.02 minutes), which is consistent with free usage as a preview and paid usage as production. The number is reported as context for the cohort definition and is not used as a finding about behavior.

Within the paid cohort, minutes are concentrated. The top 10% of paying creators produced 82.5% of paid minutes; the top 1% produced 35.0%; the bottom half produced 1.3%. Enterprise contracts (111 users) account for 23.5% of paid minutes. Concentration does not mean dependence on a few accounts: the largest single workspace is 2.3% of paid minutes and the five largest together are 8.6%.

Paid vs free, same window and exclusions	Paid	Free
Projects	66,846	67,307
Creators	2,902	65,568
Minutes per project (mean / median)	4.72 / 2.47	0.67 / 0.70
Projects of 5 minutes or longer (among projects with a duration)	22.6%	0.0%
Lip-sync enabled	20.3%	0.1%
Sourced from a YouTube URL	12.4%	6.9%
Target languages per creator (mean)	1.78	1.00
Creators using 2+ target languages	31.9%	0.3%
Projects per creator (mean / median)	23.0 / 4	1.0 / 1

Multi-language is a paid behavior, and it scales with the plan. Among paying creators, 31.9% dub into two or more target languages; among free creators, 0.3% do. The share rises with every tier: Starter 20.3%, Creator 38.8%, PRO 49.1%, Enterprise 63.1% (the Team tier, n = 14, is omitted for size; its value is 64.3%). The same ladder shows in mean target languages per creator: 1.29, 1.84, 2.70 and 3.11. The June edition presented the multi-language gap as its third finding. On paid data the gap persists and is associated with plan tier: it runs between plan tiers rather than between hobbyists and professionals, and it widens at each step.

The frontier is still far from the median. 172 paying creators dub into five or more languages, 44 into ten or more, and the median paying creator dubs into one. Tier here is the highest plan a creator reached in the window, so each creator counts once.

Sources: `paid-headline-2026.csv`, `paid-vs-free-profile-2026.csv`, `paid-concentration-2026.csv`, `paid-multi-language-adoption-2026.csv`, `paid-multi-language-by-tier-2026.csv`, `sensitivity-key-findings-2026.csv`.

M3 · Where the minutes go: industry

The June Use Case Map ranked industries by project count. Counted in paid minutes, the order changes and the picture gets sharper: long-form verticals rise, short-form verticals fall.

Industry (paid, Dec 2025 to Aug 2026)	Paid projects	Dubbed minutes	Share of paid minutes	Minutes per project	Lip-sync	Enterprise share of minutes
Education	**8,147**	**56,977**	**18.8%**	**7.09**	**23.7%**	**48.7%**
Religion	**6,823**	**34,320**	**11.3%**	**5.04**	**39.9%**	**10.0%**
Medical & Health	**5,388**	**29,770**	**9.8%**	**5.58**	**36.0%**	**50.1%**
Gaming	**3,718**	**28,552**	**9.4%**	**7.77**	**8.2%**	**27.3%**
Science & Tech	**4,027**	**21,889**	**7.2%**	**5.50**	**14.8%**	**18.8%**
Business & Finance	**3,344**	**14,187**	**4.7%**	**4.35**	**31.0%**	**5.1%**
News	**1,942**	**12,072**	**4.0%**	**6.50**	**14.9%**	**4.1%**

Education is the largest paid vertical by minutes, and half of it is contract volume. Education carries 18.8% of paid minutes with 7.1 minutes per project, the second-longest average among named verticals after Gaming. 48.7% of Education minutes come from Enterprise contracts, second only to Medical & Health (50.1%). Education's target mix by minutes is English 43.7%, Portuguese 9.4%, Spanish 7.7%.

Religion and Medical are the lip-sync verticals. Lip-sync is enabled on 39.9% of paid Religion projects and 36.0% of Medical & Health projects, against 20.3% across the paid cohort and 8.2% in Gaming. Both verticals put a speaking person on screen. Their plan mix differs: Religion is self-serve (76.7% of its minutes on PRO plans), while half of Medical & Health minutes come from Enterprise contracts (50.1%), so lip-sync use does not follow the contract split. Religion's target mix by minutes is English 18.1%, Spanish 17.1%, Chinese 12.2%.

Gaming is the long-form vertical. Gaming projects average 7.77 minutes, the longest of any industry, with 27.3% of minutes on Enterprise contracts. Its target mix is English 25.9%, Spanish 13.1%, Russian 8.8%.

Short-form verticals that were among the larger verticals by project count (Product Review, Lifestyle) average under 2.2 minutes per paid project and together hold under 4% of paid minutes. Counting minutes rather than projects is what moves them down.

Cells in the industry-by-language map are flagged in the long-format file where they hold fewer than 300 dubbed minutes, and flagged cells are not used for findings.

Sources: `paid-industry-minutes-2026.csv`, `paid-tier-mix-by-industry-2026.csv`, `paid-usecase-map-minutes-2026.csv`, `paid-production-signals-by-industry-2026.csv`.

M4 · Target-language economics

English is the destination for 31.0% of paid minutes. Spanish (11.8%), Portuguese (10.1%) and French (6.8%) follow, and the top four languages together take 59.8% of paid minutes. The 2026 language expansion took the number of target languages seen in use across all platform projects to 68, up from the June edition's 34. Between April 29 and August 31, paying creators used 55 of those 68; 21 of the 55 were new since June, and they account for 0.45% of the 95,458 paid minutes in that period. Breadth expanded; paid-plan usage stayed concentrated.

Target language (paid)	Paid projects	Dubbed minutes	Share of paid minutes	Minutes per project	Projects of 20+ minutes
English	19,557	94,109	31.0%	5.30	7.3%
Spanish	8,890	35,896	11.8%	4.17	3.6%
Portuguese	6,111	30,622	10.1%	5.10	4.9%
French	5,849	20,678	6.8%	3.57	2.3%
Korean	2,756	12,981	4.3%	4.91	4.7%
Chinese	2,170	12,256	4.0%	5.79	10.0%
Japanese	2,580	11,357	3.7%	4.46	5.3%
German	2,368	9,164	3.0%	3.91	3.8%
Italian	2,287	8,489	2.8%	3.74	2.5%
Russian	1,717	8,460	2.8%	5.01	3.7%
Turkish	653	7,155	2.4%	11.20	17.1%

Some languages attract long content. Turkish is the clearest case: 653 paid projects but 7,155 minutes, 11.2 minutes per project, and 17.1% of projects at 20 minutes or longer. 43.8% of Turkish-target minutes are Medical & Health. Arabic (756 projects, 6.71 minutes per project) and Chinese (5.79 minutes, 10.0% of projects at 20 minutes or longer, led by Religion at 34.2% of Chinese-target minutes) are the next longest among languages above the publication guardrail. French, by contrast, is a short-form destination: 3.57 minutes per project and 2.3% at 20 minutes or longer.

A note on labels: with the April 29 export the platform's labels for three languages switched to regional variants (English (US), Spanish (Mexico), Portuguese (Brazil)). This update collapses them to the base language throughout so that the June and September figures compare; the raw label counts are listed in `paid-language-minutes-2026.csv`.

Contract flows and self-serve flows are different maps. Measured across all paid tiers, the largest single source-to-target flow is Korean to English (50,892 minutes, 16.8% of paid minutes). 76.6% of those minutes come from Enterprise contracts. Among self-serve plans only, the leading flows are English to Portuguese (19,957 minutes), English to Spanish, English to French and Portuguese to Spanish, and Korean to English ranks fifth. Contract flows of long-form catalogue content and self-serve flows from individual creators are two different patterns of use, and this update reports them separately rather than as one trend.

Sources: `paid-language-minutes-2026.csv`, `paid-new-languages-2026.csv`, `paid-language-pairs-2026.csv`, `paid-language-pairs-by-tier-2026.csv`, `paid-largest-flow-composition-2026.csv`, `sensitivity-key-findings-2026.csv`.

M5 · Who pays, and from where

Enterprise contracts are invoiced outside the payment platform and carry no country, so this chapter covers self-serve plans only. Self-serve plans hold 59,638 paid projects and 231,910 minutes, 76.5% of paid minutes. A billing country is available for 90.7% of self-serve paying creators through their payment record, and the country tables cover that matched subset: 2,542 creators, 53,295 projects and 211,381 minutes, which is 69.7% of all paid minutes. Countries are shown where at least 10 creators and 300 minutes are present (35 countries); the remaining 72 countries are pooled, and the unmatched self-serve and enterprise residuals are explicit rows.

Billing country (self-serve paid)	Creators	Dubbed minutes	Share of matched self-serve minutes	Share of all paid minutes	Top target languages by minutes	Top industry
Brazil	422	50,854	24.1%	16.8%	Portuguese 42.5%, Spanish 19.4%, English 12.2%	Religion 26.6%
South Korea	316	25,198	11.9%	8.3%	English 35.4%, Korean 29.7%, Chinese 14.6%	Education 17.7%
United States	361	18,862	8.9%	6.2%	Spanish 22.6%, English 14.0%, Chinese 6.5%	Product Review 17.6%
France	84	10,901	5.2%	3.6%	French 82.9%, English 10.8%, Chinese 4.4%	Education 27.5%
Canada	58	8,796	4.2%	2.9%	English 41.8%, Spanish 26.0%, Russian 7.8%	Gaming 67.4%
Germany	101	8,560	4.1%	2.8%	Russian 21.5%, German 16.9%, English 15.4%	Gaming 29.1%
Other 29 countries above the guardrail					see `paid-country-profile-2026.csv`	
72 countries below the guardrail (pooled)	289	21,151	10.0%	7.0%		
Self-serve, country unmatched	261	20,529		6.8%		
Enterprise contracts, country not available	111	71,331		23.5%		

Brazil is the largest self-serve billing country by minutes. Creators billed in Brazil produced 24.1% of matched self-serve minutes, and 57.5% of those minutes target languages other than Portuguese, led by Spanish (19.4%) and English (12.2%). Religion is the leading vertical among them (26.6% of minutes). Where the dubbed videos are distributed is not observed.

Target mixes differ sharply by billing country. Creators billed in France direct 82.9% of their dubbed minutes into French, and creators billed in Thailand 75.0% into Thai. Creators billed in the United States direct more minutes into Spanish (22.6%) than into English (14.0%), and Product Review is their top vertical. The data records the billing country and the target language, not the source market of the content or the audience it reaches, so these rows describe what was produced and do not measure import or export demand.

Language breadth per creator is similar across the largest billing countries. Mean target languages per creator run from 1.4 to 2.0 among the six largest self-serve billing countries. None of them stands out as unusually multilingual once the average is taken per creator rather than per project.

Billing country is where the payment method is registered; it does not locate the audience. Country shares of all paid minutes are lower bounds, because 30.3% of paid minutes have no country.

Sources: `paid-country-profile-2026.csv`, `paid-country-language-long-2026.csv`.

M6 · Revisiting the June findings

The June edition promoted three findings. Re-run on the paid cohort with the same definitions and confidence-interval convention, one is confirmed, one is retired, and one is corrected. All three are kept in the record.

Finding 1 (June): "Religion has a dual hub, English ≈ Portuguese." Corrected.

On the cumulative categorized data the dual hub is still visible: within Religion, English 24.3% and Portuguese 21.7% (n = 12,174, CI ±0.8 and ±0.7 points). On paid Religion projects it disappears: English 27.8% (±1.1), Spanish 14.8% (±0.8), Portuguese 4.6% (±0.5), n = 6,823.

The difference sits in the free tier, as far as the data can show. In the free cohort of this update (December 2025 to August 2026), 1,467 free accounts created exactly one Portuguese-target Religion project each and almost nothing else: across the whole window these accounts averaged 1.27 projects, a mean lifted by a single account with 352. The June window itself (October 2025 to April 2026) can be split by plan only from December 1, 2025, when the credit ledger begins. In that part of the window, 924 of 1,632 free-tier Religion projects (56.6%) targeted Portuguese, each from an account with no other Portuguese-target Religion project (922 of those 924 accounts created nothing else at all), against 8.6% of the 2,529 paid-tier Religion projects; the 1,975 Religion projects of October and November 2025 predate the ledger and cannot be split. The June figure was computed on all projects together, so this update cannot state its cause for those two months; for the five months it can test (December 2025 to April 2026), the Portuguese share sits in the free tier. Paying creators in the Religion category localize into English, Spanish and Chinese. The paid figures hold when inherited-tier projects are excluded (English 35.8%, Portuguese 8.5%) and when the transition month is excluded (27.5% and 4.2%).

Finding 2 (June): "Korean is the structural #2 target in Science & Tech." Retired.

On the cumulative categorized data the June ranking is a tie within the interval: Korean 9.8% and Spanish 9.6% (n = 9,282, CI ±0.6 each). On paid Science & Tech projects, Korean is the fifth target language by project count (7.0%, ±0.8), behind English (30.2%), Spanish (10.4%), French and Italian, and sixth by minutes. Within Korean-target paid projects, Education (17.3%) and Science & Tech (10.2%) remain the leading verticals, so the narrower observation that knowledge content dominates inbound localization into that language still holds. The ranking claim does not. The result supports the June caution that the pattern might reflect user acquisition rather than demand, although single-platform data cannot separate the two. This update does not replace the finding with another single-language story.

Finding 3 (June): "The multi-language adoption gap." Confirmed: the gap persists and is associated with plan tier.

Cumulative professional creators still show a heavy tail: median one target language, mean 2.39, 476 creators at five or more, 148 at ten or more, maximum 33. The top 1% by language count (n = 43, one creator per workspace) average 21.3 target languages. Chapter M2 adds the plan-tier view: the share of creators using two or more languages rises across the plan ladder when creators are grouped by the highest plan reached in the window, from 20.3% on Starter to 63.1% on Enterprise (Team, n = 14, omitted as in Chapter M2). This is an association in a cross-section; the data does not show that upgrading causes multi-language use, or that multi-language use drives revenue.

A note on the June figure: the June edition stated that its top 1% cohort (n = 47) averaged 15 target languages. In the published June data file, 15 is the minimum of that cohort and 20.3 is the average. The September figure of 21.3 (top 1% by target-language count, one creator per workspace) is therefore close to flat against June.

Sources: `v1-findings-retest-2026.csv` (including the free-cohort Religion rows and the full paid Science & Tech target ranking), `multi-language-adoption-2026.csv`, `top1pct-cohort-anonymized-2026.csv`, `sensitivity-key-findings-2026.csv`.

M7 · Methodology and limitations, in brief

Full methodology, definitions and the change log are in `data/methodology-2026.md` (v2.0). The points a reader needs before citing this update:

Two platform changes that shape the data

First, Perso Dubbing moved from unlimited plans to usage-based pricing on December 1, 2025, and the credit ledger that records which plan paid for each project starts on the same day. Project counts before and after that date are not the same unit: an unlimited plan invites experiments, a usage-based plan does not. Monthly output comparisons across that date do not hold (the monthly series is published in `platform-monthly-minutes-2026.csv` and discussed in the methodology, Section 6.3), and this update reports no growth rate across it.

Second, on June 11, 2026 the platform stopped enabling the share link by default. The daily share-link-enabled rate was 97.3% on June 10, 57.8% on the day of the change, and 12.4% from June 12 to August 31 (13.0% for the rest of June, 13.3% in July, 11.1% in August; `share-rate-daily-2026.csv`). The June edition called this metric "share rate" and read its 96% level as a behavioral fingerprint of dubbing as a distribution act. This update does not carry that reading forward. The metric records whether a project had a share link enabled, and before June 11 that was the default state of every project; the data does not show whether the creator opened, sent or published the link. The figure is kept as a descriptive statistic on both sides of the change and published in full: 95.9% of projects through June 10, 2026, and 12.4% from June 12 to August 31. Neither figure is used as a finding, and the two are not comparable, because the second measures the new default rather than a choice.

- **Two sources, one pipeline.** Project rows come from two complete platform exports (January 1, 2025 to April 28, 2026, and April 29 to August 31, 2026). Plan tier and credits come from the platform's credit ledger; billing country from the payment platform; abuse and internal-account flags from platform records. Every published file is produced by a deterministic script chain; running it twice gives byte-identical files.
- **June reproduced.** Re-running the June period through the v2 pipeline reproduces 316,856 projects, 112,797 categorized projects, 909 language pairs and the 95.8% share-link-enabled rate exactly. Three June creator counts (4,023 professional creators, 115,439 single-use creators, 143,805 professional projects) do not reproduce under the documented rule (3,743, 114,573 and 145,596); all three derive from one per-account count, so an undocumented June filter is the likely explanation; the June pipeline was not retained and the exact cause cannot be recovered. This update states its rule and the recomputed cumulative figures.
- **Paid cohort rules.** Plan tier is taken from the credit ledger at the time of use. Lip-sync and multi-language child projects are charged on the parent and inherit the creator's tier for that month; they are 35.5% of paid projects and 26.5% of paid minutes, and every headline is also reported without them. Free-tier projects, multi-account abusers, internal and demo workspaces, and cancelled projects are excluded.
- **Minutes.** Dubbed minutes are the translated-video length. 96.0% of paid projects carry a duration; audio-separation and transcription projects do not, and averages use only projects with a duration.
- **December 2025 is a transition month.** Legacy balances were consumed after the pricing change; the month holds 30.3% of paid minutes in the window. All headline figures are reported with and without it.
- **What this data cannot say.** It covers Perso Dubbing's paying customers and is not an estimate for the AI dubbing market. Billing country is neither the audience country nor the source market of the content. Contract volume is concentrated in a small number of accounts and is reported separately from self-serve behavior. No year-over-year growth rate is reported.

Three things to take away from the mid-year update

1. **Measure paid usage in minutes.** Once trial output is removed, paid-plan dubbing usage on Perso Dubbing is concentrated, long-form and vertical: Education, Religion, Medical and Gaming hold 49.3% of paid minutes, and the top 10% of paying creators produce 82.5% of them. Within this data, minutes per vertical describe paid-plan usage better than project counts do.
2. **The multi-language gap is a plan-tier gap.** Grouping each paying creator by the highest plan reached in the window, two-or-more-language adoption rises from 20.3% on Starter to 63.1% on Enterprise (Team, n = 14, omitted for size). The gap the June edition identified persists, and it is associated with plan tier. This is a cross-section: higher tiers permit more languages, and creators who need more languages choose higher tiers. The data does not separate the two, and it does not show that upgrading causes multi-language use.
3. **Free usage and paid usage are different maps.** Within Religion, Portuguese takes 4.6% of paid projects (December 2025 to August 2026, n = 6,823) against 21.7% of all categorized projects (October 2025 to August 2026, n = 12,174). In the free cohort of the same paid window, 1,467 accounts created one Portuguese-target Religion project each. Within Perso Dubbing's data, a figure that mixes free and paid usage describes a different population from one that counts paid usage alone, and readers should check which of the two a figure describes before comparing it with anything else.

Full methodology, definitions and change log: data/methodology-2026.md (v2.0). Every figure above appears in a public data file; the file is named under each chapter.

 · Cite as: Kwon, T., and Perso Dubbing Data Team. (2026). State of AI Dubbing 2026: A Multi-Vertical Analysis, Mid-Year Update (v2.0). https://perso.ai/research/state-of-ai-dubbing-2026/