

THE STATE OF ATTENTION 2026

One account is no longer enough.

The brands winning short-form aren't posting better content. They have fleets of accounts, and only one brand in a hundred does.

27x

the views of the typical fleet brand,
against the typical brand posting from
one account

64%

of all views go to the top 1% of posts

1.4x

viral odds when the product is shown in use

5T+ views · 49M+ social posts in Maya's dataset · 440B views analysed for this edition

Research by Maya, the AI CMO, with Claude · Foreword by Mitul Bharti · An Attention Arbitrage publication

The report on one page

Attention is a portfolio problem.

A handful of posts take almost all the views. Nobody can predict which ones. Accounts get better as they grow. So the brands winning short-form don't bet on one account. They have fleets of them.

49M+ social posts in Maya's dataset, 5T+ views	4.2M posts analysed for this edition, 440B views	10,000+ brands and publishers behind them	3,199 brands that sell, with 50+ posts each	32 of them have a fleet
--	--	---	---	-----------------------------------

Of those 3,199 brands with 50 or more posts, 32 have five or more accounts attributed to them. That is one brand in a hundred. On average a fleet collects 20x the views of a brand posting from one account; at the median, the typical fleet collects 27x. The gap is not a secret format: a fleet post breaks out only a little more often than a brand-account post, 14.2% against 12.2%. The 20x lives in the total. More accounts, more attempts, more time in market. Nor does it take a famous brand: networks of small accounts have taken Goodnotes, Glam AI and Cantina past 300 million views each.

	Brand account	Paid creators	Fleet
Views per brand, average	34M	143M	696M
Views per brand, median	5.0M	50.5M	136M
Views per post, median	36K	30K	93K
Breakout rate	12.2%	11.1%	14.2%

Brands with 50 or more posts. Tracked views, not unique reach or sales. Full table in Part 2.

Three facts the market is ignoring

Concentration. The top 1% of posts take 64% of the views, and the bottom half take a third of one percent. Optimising the average is the wrong job.

Unpredictability. The median account breaks out, at 5x its own median views, on one post in nine, and no creative choice we tested moves those odds by more than about 1.6x. You can't pick the winner. You can only make more real attempts.

Compounding. Breakout rates rise from 5.9% for accounts under 1,000 followers to 17.9% for accounts over a million. Hold more accounts, for longer, and the portfolio gets better.

What moves a post: same creator, not pooled folklore

Show the product in use: 1.4x the odds of going viral. Go long: one-to-three-minute videos break out 1.42x as often as 8-to-15-second clips. Write a real caption. Hook templates move the odds about 10% either way. Testimonials, lo-fi footage, selfie video and the posting day all look like they work in pooled data, and do nothing once the creator is held constant.

Fully AI-made video looks 56% worse when you pool every account. Within the same creator it performs on par with filmed video. Volume is what kept fleets at one brand in a hundred, and that constraint is coming down.

Contents

BEFORE YOU START

The report on one page	2
Foreword	4

THE FINDINGS

Part 1	Three facts about attention	5
	Concentration, unpredictability, and accounts that compound	
Part 2	Three ways to run a brand on short-form	7
	Brand accounts, paid creators and fleets, compared	
Part 3	Inside the fleets	9
	Portfolio fleets, network fleets, and what they share	
Part 4	The Creator Illusion	10
	Why the popular short-form rules fail within the same creator	
Part 5	What moves a post	11
	Hooks, product visibility, length, format and captions	
Part 6	AI loosens the constraint	13
	AI-made versus filmed video, measured within the same creator	
Part 7	Operating implications	14
	For brands, creators, agencies, investors and boards	
Part 8	Five calls for the next two years	15

REFERENCE

The playbook	16
Method	17
Benchmarks	18
About Maya	20

Foreword

Anyone can build now. Almost no one can market.

Software got cheap. A small team can ship in a weekend what used to take a year. What didn't get cheap is attention: getting the right people to stop, watch and buy. That is the CMO's job, and most teams hit that wall long before they can afford the department. So founders become the CMO at night, and the marketing gets whatever is left of them.

Maya is the CMO they hire instead. You don't operate her like software. You give her an objective, a budget and the context, and she runs the loop: decides what to make, makes it, ships it across every account the brand runs, reads what lands, and brings the calls that need a human back to the team. Then she does it again, every day, on every account at once.

A CMO is only as good as her read on the market, so Maya studies the feed relentlessly: which posts break out, from which accounts, for which products, and why. That record now spans more than 49 million social posts and more than 5 trillion views, over 3 trillion of them organic views on TikTok and Instagram. This report is built on a 440-billion-view cut of the organic side, covering more than 10,000 brands and publishers. We haven't seen anything like it published. Most "state of" reports in this space are a survey of a few hundred marketers. This one is what the market actually did.

We could have kept it as a private edge. We're publishing it because the industry is running on advice the data says is wrong, and the approach that clearly works is used by almost nobody. When that many people are wrong at the same time, the fastest way to win is to be the one who shows them.

The short version: attention is a portfolio problem. A handful of posts take almost all the views, nobody can predict which ones, and accounts get better as they grow. The brands that have worked this out don't bet on one account. They win through fleets of them, and on average they're pulling 20x the reach of brands posting from one. Until now a fleet needed a production operation most companies will never have. That constraint is breaking, and Maya was built for what comes after it.

This edition covers organic; paid gets its own. It is Maya's first public work, drawn from the same research she runs for every brand she works with. Take the playbook and run it yourself. Or hire the one who wrote it.

Mitul Bharti

Founder and CEO, Attention Arbitrage, Inc.

Part 1 • Three facts about attention

Before the operating models, three properties of the medium itself. Together they explain everything that follows.

Fact one: almost all the attention goes to almost none of the posts

Posts, ranked by views	Share of all views
Top 0.1%	29.4%
Top 1%	63.9%
Top 10%	93.7%
Bottom 50%	0.35%

Figure 1. How views are distributed across every post from brands that sell. Across all 440B views in this edition, publishers included, the top 1% take 65%. The shape holds in both.

One post in a thousand carries almost a third of all views. The bottom half of everything published, combined, carries a third of one percent. The median post from a brand that sells lands at about 3,200 views.

So the average post matters far less than it seems. Total reach depends mostly on how often a brand lands a post in the top percentile. Most teams are still optimising the average.

Fact two: nobody can say in advance which post that will be

Raw views mostly measure the size of the account, so to compare posts fairly we score each one against its own account. The **Breakout Score** is a post's views divided by the median views of the account that published it. A **breakout** is a score of 5 or more with at least 10,000 views. A **viral** post is a score of 100 or more with at least 100,000 views.

Measure	Value
Posts that break out	12.0%
Posts that go viral	1.8%
Breakout frequency, median account	1 in 9 posts

Figure 2. Breakout frequency. Median account: accounts with 20 or more scored posts.

The typical account clears its own baseline by 5x on one post in nine, and by 100x on fewer than one in fifty. Every creative choice we tested in Parts 4 to 6 moves those odds by somewhere between a few percent and about 1.6x. None comes close to making a breakout predictable. Short-form is a game of attempts, and an account can only make so many.

Fact three: accounts that grow get better at growing

Followers	Breakout rate
Under 1K	5.9%
1K to 10K	10.6%
10K to 100K	15.0%
100K to 1M	17.5%
1M+	17.9%

Figure 3. Breakout rate by account size, TikTok.

The Breakout Score already normalises for size, so in principle a big account should beat its own median no more often than a small one. In practice the million-follower accounts do it three times as often as accounts under a thousand, and most of that gain arrives between 1,000 and 100,000 followers. Some of the gap at the very bottom is mechanical, because a breakout needs 10,000 views and a 400-follower account rarely gets there at any multiple. Above that floor, the pattern suggests reach compounds.

Put those together and the fleet stops looking like a hack. If a few posts take everything, you can't predict which, and accounts improve with age, then one account is one ticket in a lottery with a fat tail. Ten accounts are ten tickets, and each one gets better as it grows. That's the whole argument. The rest of the report is the evidence.

Part 2 • Three ways to run a brand on short-form

Among the brands in this edition with at least 50 posts, short-form is organised in one of three ways. The **brand account** model: a brand posts from its own official handles, usually one to three of them. The **paid-creator** model: four or more real people are paid to post on the brand's behalf. The **fleet** model: five or more accounts, official or otherwise, attributed to the brand.

In this edition's data, 32 of 3,199 brands have a fleet attributed to them. Our coverage is thinner for a brand's extra accounts than for its main handles, so that count likely understates how many brands have fleets. What the data does show clearly is how much better fleets perform.

	Brand account	Paid creators	Fleet
Views per post, average	155K	60K	225K
Views per post, median	36K	30K	93K
Breakout rate	12.2%	11.1%	14.2%
Views per brand, average	34M	143M	696M
Views per brand, median	5.0M	50.5M	136M

Figure 4. Performance by operating model, brands with 50 or more posts. Averages are pulled up by a few very large brands, so medians show the typical brand in each model.

Fleets win on every row, on averages and on medians. On averages, a fleet brand collects 4.9x what a paid-creator brand collects and 20x what a brand account collects. Averages flatter the biggest brands, so look at the typical brand too: the median fleet has 136 million views against 5 million for the median brand account, a 27x gap, and 50.5 million for the median paid-creator brand, a gap that narrows to 2.7x. Remove Duolingo, the largest fleet in the data, and the remaining 31 still average 442 million: 3.1x the paid-creator figure and 13x the brand-account figure. No single outlier is doing the work.

The paid-creator model

Paid creators are a common answer to “we need more short-form”, and on reach per post they are the weakest of the three models in our data. Flo Health illustrates the pattern: of the 40-plus accounts in the data for the brand, nearly all belong to creators rather than to Flo, and the programme produces a great deal of output at modest reach per post. Sixty thousand views per post is not nothing. It is a third of what a fleet returns for the same post, and the ratio holds at the median: 30K against 93K.

A fair caveat: we measure views, not cost and not conversion. A creator programme can still be the right call for a brand that is buying trust or a specific audience. But when a programme is bought for reach, it is buying the weakest-performing model of the three in our data.

When one account is enough

Some brands do very well from a single handle, and they share a profile: an existing audience, or a product that sells itself on camera. The brand accounts that break out most often, Airtasker, Fabulous, BBVA, Instax, Nykaa, KineMaster, Beatport and Falabella, do so on roughly a third of their posts, close to 3x the market rate.

Few do it alone: many of their biggest posts are promos and joint campaigns with creators, and Maya credits a post's full views to the account she sees it on, collab or not. If you already have a fanbase or a product people cannot stop filming, the single account works. If you have neither, the numbers say you are in the wrong model.

Model	Brand	Category	Accounts (at least)	Views	Breakout rate
Fleet	Grab	Consumer services	5+	5.1B+	30 to 40%
	Duolingo	Education	15+	8.5B+	10 to 20%
	ArtWorkout	Education	60+	1.0B+	10 to 20%
	Lyfta	Health & fitness	10+	1.0B+	10 to 20%
	Nike	Retail & fashion	5+	1.0B+	10 to 20%
Paid creators	ChatGPT	AI	5+	2.1B+	under 10%
	Flo Health	Health & fitness	40+	1.2B+	10 to 20%
	Bump	Social	100+	2.9B+	10 to 20%
	Equals	Social	20+	1.0B+	10 to 20%
Brand account*	OpenAI	AI	1 to 3	1.0B+	10 to 20%
	Swiggy Instamart	Consumer services	1 to 3	2.7B+	10 to 20%
	Osmo	Devices & creative tools	1 to 3	2.5B+	20 to 30%
	Instax	Devices & creative tools	1 to 3	1.2B+	30 to 40%
	King English Kids	Education	1 to 3	1.3B+	10 to 20%
	BBVA	Finance	1 to 3	1.1B+	30 to 40%
	Mercado Pago	Finance	1 to 3	1.8B+	10 to 20%
	Nykaa	Retail & fashion	1 to 3	2.3B+	30 to 40%
	Palmonas	Retail & fashion	1 to 3	1.0B+	20 to 30%

Figure 5. Brands with a product to sell that passed one billion views, September 2026, grouped by operating model, then category. Accounts and views are banded lower bounds for the accounts in Maya's data, not a complete inventory.

* Maya credits a post's full views to the account she sees it on, including creator collabs. A post seen on more than one account counts toward each, so totals can overlap.

The yardstick

The publishers sit outside this analysis, but one comparison is worth making. The flagship accounts in the data for Disney (1.5B+ views), UFC (1.2B+), ESPN (1.2B+) and the NBA (1.0B+) each clear a billion on their own. Duolingo's 15-plus accounts have collected more than 8.5B between them, more than those four handles combined. Each publisher runs many more handles than its flagship, so read it as scale, not a league table. But scale is the point: a language app, run as a fleet, out-reaches the front doors of the biggest entertainment brands in the world.

Why fleets are rare

If fleets are this much better, why are they this rare? Because they're brutally hard to run by hand. Every account needs a constant supply of content, someone to publish it and someone watching what lands, and for most teams one account already eats that capacity. Part 6 is about what happens when volume stops being the constraint.

Part 3 • Inside the fleets

Fleets come in two shapes, and the distinction matters because only one of them is available to a brand without a household name.

Portfolio fleets are what large brands build when they split their presence by market, product line or audience. Duolingo’s 15-plus accounts have passed 8.5 billion views between them; Grab’s have passed 5 billion and Nike’s a billion. These fleets are visible, and they lean on a brand people already know.

Network fleets are the quieter form and the more useful one to study. They are assembled from many small accounts, most of which no agency would call a creator.

Fleet type	Brand	Accounts (at least)	Views	Breakout rate	AI-made	Faceless
Portfolio	Duolingo	15+	8.5B+	10 to 20%	~20%	~50%
	Grab	5+	5.1B+	30 to 40%	under 10%	~20%
	Nike	5+	1.0B+	10 to 20%	<i>not tagged</i>	<i>not tagged</i>
	Glovo	5+	370M+	20 to 30%	under 10%	~20%
	Canva	5+	80M+	10 to 20%	<i>not tagged</i>	<i>not tagged</i>
Network	Goodnotes	30+	630M+	10 to 20%	under 10%	~80%
	Glam AI	15+	510M+	10 to 20%	~70%	under 10%
	Cantina	30+	320M+	10 to 20%	~100%	~80%
	Spotify	5+	230M+	10 to 20%	under 10%	~100%
	Halo AI	15+	160M+	10 to 20%	~10%	~80%
	Roblox	5+	90M+	10 to 20%	~60%	~80%
	Wondershare Filmora	5+	60M+	10 to 20%	~40%	~50%

Figure 6. Selected fleets from Maya’s wider data, September 2026. Portfolio fleets are mostly official accounts; network fleets are mostly small creator accounts attributed to the brand. “Not tagged” means no creative read for that brand.

There’s no house style. Spotify’s network is almost entirely faceless and barely uses AI; Cantina’s is almost all AI-made. Glam AI’s accounts mostly show a face; Goodnotes’ mostly don’t. The fleets disagree on every creative decision. The one thing they have in common is structural: many accounts, each making its own attempts, all pointed at one product, with the brand collecting the breakouts wherever they land.

Cantina is the one to remember. Across the 30-plus accounts we track, the typical account has about 2,000 followers, which in most marketing meetings would get read as a failed launch. Together they have passed 300 million views, and almost all of it is AI-made. A portfolio fleet needs a famous brand. A network fleet needs an operating system. And every account in it still needs a job: a language, a use case, an audience, a point of view. Handle count alone is not a strategy.

Part 4 • The Creator Illusion

Most short-form advice gets made the same way. Someone collects a few hundred viral posts, notes what they have in common, and publishes the common features as rules. It feels empirical. It has a flaw that no sample size fixes.

Viral posts are not a random draw. They come overwhelmingly from large, established accounts, and large accounts differ from small ones in what they post, how they post, and what audience is waiting when they do. So the features that viral posts share are, to a great extent, the features that big accounts share. The method mistakes the habits of the successful for the causes of success. We call this the **Creator Illusion**, and it is the reason this report measures every creative choice within the same creator: a creator's posts that do the thing, against the same creator's posts that do not, so that audience and account size cancel out.

The table shows the same rules measured both ways.

Rule	Pooled comparison	Same creator	Verdict
Use testimonials	1.83x	1.02x	No effect
Handheld selfie footage	1.43x	0.87x	Hurts
Silent reaction with text	1.41x	0.99x	No effect
Raw, lo-fi footage	1.35x	0.98x	No effect
Problem, then solution	1.32x	0.91x	Slightly hurts
Stitches and reactions	0.84x	1.14x	Helps
Faceless content	0.77x	1.09x	Helps
Trending sounds	0.66x	1.09x	Helps

Figure 7. Relative odds of going viral, pooled across all accounts versus within the same creator. 1.00x is the creator's own baseline.

It cuts both ways. Testimonials look 83% better in the pooled view and do nothing once the creator is held constant. Raw phone footage, the “authentic” look a whole production style is built on, does nothing. Handheld selfie video is actively worse. Faceless content, stitches and trending sounds all look like losing moves in the pooled data and turn out to help. The pooled view wasn't wrong about what viral posts look like. It was wrong about why they went viral.

That has a cost. A brand that scripts testimonials, stages a lo-fi look and briefs its creators to film selfie-style is paying for three choices that, inside its own accounts, are worth nothing or less. A viral example is a starting point for a hypothesis, not an explanation of why it succeeded. A library of winners cannot tell a habit from a cause.

Even the oldest rule of all, the right day to post, comes out flat.

Day	Mon	Tue	Wed	Thu	Fri	Sat	Sun
Breakout rate	12.1%	12.1%	12.1%	12.1%	12.0%	11.6%	11.9%

Figure 8. Breakout rate by day of posting.

Part 5 • What moves a post

So what's real? Measured within the same creator, a short list of unglamorous things.

The hook is not the lever

No part of a short-form video gets more attention than its first two seconds, and there is an entire cottage industry selling hook templates. So we tested every major hook category against the same creator's other posts.

Hook type	Same creator
Bold claim / result	1.09x
Problem callout	1.06x
Pattern interrupt	1.04x
Story opener	1.03x
Social proof	1.02x
Curiosity gap	1.00x
Relatable callout	0.99x
How-to promise	0.97x
Urgency / FOMO	0.86x
Contrarian opener	0.64x

Figure 9. Hook category, relative odds of going viral within the same creator.

Eight of ten categories sit between 0.97x and 1.09x. The two outside that band, urgency and the contrarian opener, are the least used and rest on around 50 creators each, so don't lean on them. None of this means the opening doesn't matter; a flat first second still loses the viewer. It means the category of hook, which is the only thing a template can give you, is worth about 10% either way.

What actually moves it

Creative choice, same creator	Effect
Video vs carousel	1.55x
60s to 3min vs 8 to 15s	1.42x
Product shown in use vs not shown	1.39x
7+ hashtags vs 1 to 3	1.31x
Confessional, shortcut or opinion formats	1.23x to 1.31x
16 to 40 word caption vs 1 to 5 words	1.25x
Disclosed as an ad	1.05x
Listicle	0.87x
Green-screen explainer	0.74x
No caption	0.70x

Figure 10. Creative choices measured within the same creator. Format, length, caption and hashtag effects are at the breakout threshold; the rest at the viral threshold.

Showing the product in use lifts the odds of going viral by 1.4x against the same creator's posts that keep it out of frame. That's the opposite of what many creator briefs ask for, which is to hide the product so the post "doesn't feel like an ad". Posts disclosed as ads perform on par with, or a little better than, the same creator's organic posts. In these data, audiences aren't punishing brands for selling. They're rewarding brands for showing what they sell.

Length is the other one that overturns received wisdom. After years of pressure towards shorter and shorter clips, one-to-three-minute videos earn a median of 9,364 views against 1,960 for 8-to-15-second clips, get shared 1.7x as often on TikTok, and within the same creator break out 1.42x as often. In these data, the audience wants longer video. Adding seconds won't save a weak idea, but an idea that needs a demonstration, a progression or a payoff should be given the time to land.

The rest is craft. Video beats carousels by half. A caption of real length beats a five-word one by a quarter, and no caption at all costs almost a third. Formats with a point of view, the confession, the shortcut, the opinion, outperform. The green-screen explainer, a staple of the software category, underperforms by a quarter.

Part 6 • AI loosens the constraint

Part 2 ended on why fleets are rare: they need more content than a human team can make. That makes this the most important table in the report.

Content type	Pooled comparison	Same creator
Filmed	1.00x	1.00x
Partly AI-made	0.68x	1.17x
Fully AI-made	0.44x	1.06x

Figure 11. Relative odds of going viral, AI-made versus filmed. The within-creator figure for fully AI-made video rests on 110 creators who published at least three such videos.

Pool everything together and fully AI-made video looks like a disaster, 56% less likely to go viral than filmed. That’s where the near-universal belief that platforms suppress AI content comes from. Hold the creator constant and the penalty disappears: fully AI-made video performs on par with the same creator’s filmed video, and partly AI-made video does better.

The pooled penalty looks like the Creator Illusion again: because the gap closes within creators, it reflects which accounts publish AI-made video, not how that video performs. This is an observation within the accounts in this edition, not a test of any platform’s policy, but we find no trace in these data of a penalty on AI-made video itself.

The main thing that made fleets expensive was feeding them. If AI-made content holds its own against filmed content from the same account, the barrier that kept the fleet model rare is coming down.

Easier production raises the cost of bad judgment. Ten versions of the same weak assumption generate activity without generating learning. Capacity is only useful when it lets a team test a real question and act on the result.

Part 7 • Operating implications

For brands. The question is no longer “what should our account post this week” but “how many credible attempts are we making, across how many accounts, and is each of them beating its own baseline”. Score every account against its own median, not against raw views. Put the product on screen. Publish longer than feels safe. Pull the per-post reach on your creator programme before renewing it. Give each account a job before you open it.

For creators. Most of what you have been told about hooks and timing is noise. Show the product, go longer, take a position, and measure yourself against your own normal rather than someone else’s highlight reel.

For agencies. The casting-and-paying model returns the fewest views per post of anything we measured. The brands pulling away sit at the centre of account networks, and they will hire partners who can build and operate one.

For investors and boards. Follower count describes one account’s past. Account count and attempts per week describe a brand’s capacity to produce breakouts in the future. The fastest-compounding brands in this data are the ones that treated distribution as infrastructure to build, not a channel to post into.

On running a fleet properly. Every fleet in this report that works is a set of real accounts: distinct content, a point of view per account, the product in plain sight. Copy-pasting one video across forty handles isn’t a fleet, it’s spam, and platforms treat it that way. The edge comes from more real attempts, not more copies.

On what this report does not measure. Views are not unique people, visits or sales, and production cost is not in the tables. Gaps between models can also reflect brand, budget, product and observation window. Connect the distribution to your objective before you scale the operation.

Part 8 • Five calls for the next two years

These are expectations, not findings; a cross-section of 2025 and 2026 cannot show a shift in motion. Future editions will track them.

- 1 The fleet becomes the default operating model** for brands that sell on short-form, and “how many accounts” becomes a normal question in a marketing review.
- 2 Account count joins the growth dashboard**, reported to boards next to acquisition cost and retention, with the job of each account written next to the number.
- 3 Paid-creator budgets get re-underwritten** on reach per post, and a meaningful share of them moves to owned networks.
- 4 Hook templates lose their market** as teams measure within their own accounts and discover the category was never the lever.
- 5 Baseline-relative scoring becomes standard.** Raw views mostly measure the size of the account; a Breakout Score measures the post. The industry will start reporting the second.

The playbook

Do more of this	Evidence	Do less of this	Evidence
Run many accounts as one system	20x views per brand on average, 27x at the median	Post from a single brand account	34M average, 5M median views per brand
Increase credible attempts	1 in 9 posts breaks out	Bet on individual posts	Breakouts are not predictable
Show the product in use	1.4x	Rotate hook templates	Most within $\pm 10\%$
Publish video, not carousels	1.55x	Stage a lo-fi look	No effect
Go to 60s to 3min	1.42x	Chase a posting day	11.6% to 12.1% all week
Use confessional, shortcut, opinion formats	1.23x to 1.31x	Use green-screen explainers	0.74x
Write 16 to 40 word captions, 7+ hashtags	1.25x, 1.31x	Post without a caption	0.70x
Feed every account with AI-made content	No penalty	Learn from small viral samples	You are measuring the creator
Give each account a job	Networks work when accounts differ	Clone one video across handles	Spam, not a fleet

Method

The dataset and this edition. Maya's dataset spans more than 49 million social posts, organic and paid, with more than 5 trillion views between them; the organic side alone carries more than 3 trillion views across over 20 million public TikTok and Instagram posts. This edition is built on the September 2026 cut of the organic data: 442.9 billion views across 4.2 million posts from over 10,000 brands and publishers, 91% of it published in 2025 and 2026, the portion of the corpus Maya had processed when this analysis ran. It includes media, sports, news and entertainment publishers, for whom the content is the product. This report analyses the part of it where that is not true: brands that use short-form to sell a product, together accounting for over 174 billion views. Publishers are excluded from every figure except the two places the text names them, the distribution comparison in Part 1 and the yardstick in Part 2.

Coverage. Account counts, follower figures and brand totals describe the accounts and posts in this edition's data, observed up to September 2026, not a census of either platform or of what a brand operates. Figures 5 and 6 draw on Maya's wider data and show banded lower bounds. Large brands and publishers in particular run more accounts than this edition captures, brands open and close accounts continually, and the operating-model classification reflects our coverage rather than a complete inventory.

Scope of the analysis. Consumer apps, retail, beauty, fitness, food, finance, gaming, education, software and technology brands, across their official accounts, the creators posting on their behalf, and the wider account networks some of them run.

Data. Views, likes, comments and shares were recorded per post. Video content was analysed for creative attributes: hook type, format, product visibility, presence of a face, use of AI. Instagram publishes no view counts for image and carousel posts, so view-based figures cover video.

Scoring. Each post receives a Breakout Score: its views divided by the median views of the account that published it. A breakout is a score of 5 or more with at least 10,000 views; a viral post is a score of 100 or more with at least 100,000 views.

Within-creator comparison. For each attribute we compare a creator's breakout or viral rate on posts with the attribute against the same creator's rate on posts without it, using only creators with at least three posts on each side. Where a figure compares two opposing choices it is the ratio of the two rates. Content attributes (hook, product visibility, style, AI use) are measured at the viral threshold; format, length, caption and hashtag effects at the breakout threshold.

Operating models. Assigned to brands with at least 50 posts. Brand account: one to three official handles. Paid creators: four or more accounts, mostly real people posting for the brand. Fleet: five or more accounts attributed to the brand, either a portfolio of official handles or a network of mostly faceless, AI-made or slideshow accounts. Attribution does not establish who operates an account: a network fleet's accounts may be run by the brand, by contracted creators or by independent fans.

Limitations. Everything here is observational. Comparing creators with themselves removes the largest bias in this kind of work, account size and audience, but not every difference between posts. Gaps between operating models may partly reflect brand, budget or product; the network fleets, built from small accounts, argue that the advantage is not simply brand size, but the data cannot fully separate the two. Maya credits a post's full views to each account she sees it on, including creator collabs, so brand totals can overlap and brand-account totals include creator-driven reach. We measure reach and engagement, not production cost or conversion. Follower-based benchmarks use TikTok, where follower counts are consistently available.

Benchmarks

Place your own accounts against these. A post is doing well if it clears the 90th percentile for its account's size band. An account is doing well if its breakout rate beats the band's. All figures are for brands with a product to sell.

Overall	Value
Median views per post	3,153
90th percentile post	113,679 views
99th percentile post	2.5M views
Median like rate	2.6%
90th percentile share rate (TikTok)	0.62%
Breakout rate	12.0%
Viral rate	1.8%

Posts with at least one recorded view. Like and share rates for posts with 100 or more views.

TikTok, by followers	Median views	90th pct	99th pct	Like rate	Share rate	Breakout rate
Under 1K	887	6,470	229,099	3.4%	0.01%	5.9%
1K to 10K	1,277	20,400	644,756	3.2%	0.08%	10.6%
10K to 100K	2,772	65,309	1,376,022	2.6%	0.08%	15.0%
100K to 1M	11,259	277,100	4,400,000	3.0%	0.09%	17.5%
1M+	48,972	1,384,923	13,333,513	3.3%	0.10%	17.9%

By video length	Median views	Share rate (TikTok)	Breakout rate
Under 8s	1,678	0.05%	9.9%
8 to 15s	1,960	0.07%	10.6%
15 to 30s	3,723	0.08%	12.4%
30 to 60s	5,577	0.09%	13.6%
60s to 3min	9,364	0.12%	15.2%
3min+	5,338	0.08%	14.5%

Follower and length tables: posts with 100 or more views; rates are medians per post.

By category	Median views	Breakout rate
Travel & outdoors	3,319	15.2%
Home & lifestyle	8,095	13.7%
Finance	1,549	13.4%
Food & cooking	18,879	13.3%
Gaming	7,116	12.9%
Entertainment apps	4,050	12.8%
Software & apps	3,468	12.8%
Technology & AI	3,054	12.4%
Career & work	1,390	11.5%
Health & fitness	5,315	11.2%
Business	2,693	11.0%
Relationships & family	2,165	10.8%
Beauty & fashion	12,305	10.7%
E-commerce & retail	6,011	10.7%
Education	1,943	10.6%
Marketing & sales	1,134	9.2%

Scored posts from brands with a known category.

About Maya

Maya is an AI CMO: the marketing lead a team hires, not a tool it operates. Give her an objective, a budget and the context, and she runs the loop end to end. She learns what breaks out in a brand's category, makes the creative for every account the brand runs, publishes it, reads the results, and brings the decisions that matter back to the team.

Her edge is the feed: TikTok, Instagram, creators and paid social, where attention is won or lost. She was built to run the kind of fleet this report describes, the benchmarks above are the ones she holds every account to, and the research in these pages is the same research she runs for every brand she works with.

Maya is in private beta with a select group of brands. Request access at hiremaya.ai.

Who's running marketing? Maya.

hiremaya.ai

The State of Attention 2026. By Maya, the AI CMO built by Attention Arbitrage, Inc., with Claude, an AI model by Anthropic. Foreword by Mitul Bharti. Published by Attention Arbitrage, Inc.

Disclaimer. This report was produced by AI end to end. Maya assembled and analysed the data, and Claude co-wrote the text and checked the figures for internal consistency. No person audited every number, and some details may be wrong. Attention Arbitrage, Inc. publishes this report and has a commercial interest in the practices it discusses. Treat the figures as directional, and tell us what we got wrong at hiremaya.ai so the next edition can fix it.

Maya assembled and analysed the dataset. Claude co-wrote the report: building the argument, checking the figures for internal consistency, and drafting the text. Figures are as measured by Maya, September 2026.

Citation: Maya and Claude (2026). "One account is no longer enough." The State of Attention 2026. Attention Arbitrage, Inc.