

PAID SOCIAL / STRATEGY

Meta account structures that survive algorithm shifts

A practical account operating system for Andromeda-era growth.

Campaign maps, delivery signals, retargeting, offer tests and a disciplined testing cadence.

THE PLAYBOOK IN ONE LINE

DIAGNOSE. SEGMENT. TEST. REPEAT.

A practical guide for DTC founders, operators and growth teams.

The executive read.

What matters, before you touch a budget.

01

Read the account differently

Track CPMR, frequency, 2-3 day results and the true new-customer outcome.

02

Choose ABO or CBO by evidence

Use forced learning where necessary; use CBO to scale when it works.

03

Protect prospecting intent

Separate retargeting and exclude existing customers where appropriate.

04

Budget for discovery

Set aside 10-20% for structured tests across five performance levers.

THE FIELD GUIDE

01 Who this is for

02 What Andromeda changed

03 The account map

04 Reading the account

05 Retargeting after Advantage+

06 It is always the offer

07 The circle, and the 10-to-20% rule

The following pages retain the authors' working recommendations and case examples. Diagrams and charts are editorial visualizations of the figures in the supplied drafts, not independent performance benchmarks.

Who this is for

This is the Meta account structure DigiCom builds for every brand it takes on, written for the operator spending \$25,000 or more a month who watched a working account fall off a cliff sometime in the last year and could not say why.

Andromeda, Meta's creative-retrieval model, changed how ads enter the auction, how often the same people see them, and how fast a winner stops working. Most of the playbook brands still run was written before it. This paper is the version we run now, across 200+ brands and \$200M+ in managed spend, with the numbers we use to make each call.

It covers five things: what actually changed, the campaign and ad-set map we use for accounts from \$10,000 to \$20,000+ a day, how we read an account in a 2-to-3-day window instead of day over day, how we retarget now that Advantage+ wants to do it for us, and why, when all of that is right and the account still will not scale, it is the offer.

By Hemant Varshney, Danielle Dror and Mike Constandatos, co-founders of DigiCom, with Jake Orenstein, Director of Growth. Drawn from our Performance Anxiety and Chew on This sessions through 2026.

What Andromeda changed

Four things changed, and every structure decision below exists to deal with one of them.

Fewer of your ads ever enter the auction. Andromeda scans every creative you upload and groups the ones that look alike under a single entity ID: same creator, same setting, same text-overlay style, even two cuts of the same video. Meta then treats that group as one ad. That is why accounts now show tens of ads spending a few cents each. They are not losing the auction; they are never being entered in it, because a sibling already was.

Frequency rose on consolidated accounts. If you run one consolidated setup with prospecting and retargeting in the same campaign, the algorithm shifted hard toward engaging people who already know you. Frequencies climbed to levels we had not seen before, with CPMs rising at the same time. The algorithm took the easy way out, and a metric almost nobody tracked a year ago, CPMR (cost per thousand reached, not served), became the first thing we look at.

Day-over-day reads stopped meaning anything. The same creative that crushes on Monday can do nothing on Tuesday and come back on Wednesday with three latent conversions. Judging an ad on a single day now produces the wrong call more often than the right one.

Relevance started paying. The one shift in Andromeda that works in your favor: the model rewards continuity between the ad, the person in it and the page it lands on. A supplement brand we manage had creatives that would not spend at all. The only change was putting the doctor from the ad at the top of the landing page. The same ads scaled 5 to 6x week over week at \$20,000+ a day, not because the page converted better than the PDP but because CPMs and CPCs dropped into a new audience pocket.

Not every account has every problem. Meta rolls Andromeda features out unevenly, so two accounts in the same vertical can behave differently in the same week. Which is why the first step in every audit is to diagnose which of

the four you actually have before touching creative.

The account map

Protect your next winner: 10-20% test budget

10%-20%

of monthly media budget

held for experiments, every month.



20% experimentation illustration

80% other spend

Illustration of DigiCom's 10-20% planning rule; actual allocations vary by account.

Use the budget tool that is working in your account, not the one that won the last argument on LinkedIn. We run both, and the split is decided by what the account is doing, not by doctrine.

Setup	What it is for	When we use it	What it costs you
ABO (ad-set budgets)	Fast test-and-learn. Forcing spend onto a batch so every ad gets a read.	New concepts, new personas, accounts where CBO locks onto one ad and starves the rest	Needs creative volume: in four days you can run out of ads if nothing picks up
CBO (campaign budget)	True batch tests and scale. Meta rewards it with lower CPMs because you are playing by its rules.	Mature accounts with even spend across ad sets, accounts where CPMs sit well above the vertical benchmark	One ad takes most of the budget; the rest never get a fair read
CBO leftovers into ABO	Take the ads CBO would not spend on, put them together in one ABO ad set, force the test	Accounts running primarily CBO	Works a handful of times; if Meta still will not spend, move on

Two rules sit on top of that table. First, if you are spending under \$15,000 to \$20,000 a day there is no reason to run a traffic campaign. The algorithm will send you people most likely to click, not most likely to buy, and for CPG, pet, wellness or travel there are already enough people in market to run conversion campaigns only. Second, when CPMs are far above where the vertical should be (you should know that number; \$90 to \$100+ in some categories is real, \$45 in a \$20 category is a structure problem), try the campaign style with lower CPMs before you touch

creative. Usually that is CBO.

The orchestra ad set. Because of entity IDs, an ad set full of look-alike creatives is one ad in Meta's eyes and only one of them spends. We launch ad sets the way you fill an orchestra: two videos, three carousels, a GIF, three statics, eight to ten ads that cover every format and every stage of the funnel. Spend distributes far more evenly, you get every placement Meta offers, and you are doing what Meta wants, which it rewards.

Split by margin profile. A 15% margin product and a 70% margin product have different break-even ROAS, different allowable CAC and different contribution margins. Running them in one campaign means one of them is always being judged by the wrong number. We break campaigns out by margin profile or by product, depending on the account, and we re-test the structure every time growth slows, because structure is never finished.

Budget the test before the winner dies. Every account carries 10 to 20% of monthly spend for new tests, every month, no matter how well the current winner is doing. A great creative can carry an account for weeks. It will stop. The brands that get hurt are the ones that were not already testing the next thing when it did.

Reading the account

Give the ad time to tell the truth.

2-3 DAYS

Rolling read window

1.5-2x TARGET

Pause trigger (contextual)

2-4 WEEKS

Time to give a new ad

These are operating guidelines from the manuscript, not universal platform rules.

Anchor ads earn the top-of-funnel job.

ANCHOR AD



\$0.25-\$0.30

ACCOUNT AVERAGE



\$1.60

CPC comparison from one account anecdote; the range is shown at its midpoint.

Judge ads on a 2-to-3-day window. Pause and un-pause more than you used to. Never kill an anchor ad because of its CPA.

The window. Day-over-day is dead. We read creative on a rolling two or three days, and in our best accounts we pause an ad that is over target, then turn it back on when latent conversions land. After a few cycles you see which ads keep doing that and you leave them on. Seven days is too quick to call an ad. Give it two to four weeks, even sitting at a few cents of spend, before you stop waiting and move to the next creative.

The pause rule, as a moving goalpost. The right pause point depends on how much room the account has, so we move it. If the account is beating target you can let an ad run to about 2x target before pausing. If the account is tight, pause at 150 to 175% of target before a second conversion, because a \$200 CPA on one creative is what you get for waiting for the 3x you told yourself you would allow. The sequence that matters:

- The ad converts once near target. It stays on to convert again.
- It runs to 1.5 to 2x target without a second conversion. Pause it.
- Latent conversions bring it back under target. Turn it back on.
- It keeps doing that. Leave it on and stop checking it daily.

Anchor ads. Every scaling account has one or two ads that are not the most efficient but are the volume drivers: they reach new people at a low cost per click, they fill the top of the funnel, and the efficient ads behind them only convert because of the traffic they bring. The signs: CPC far below the account average (we have an ad at 25 to 30 cents in an account averaging \$1.60), frequency under 2 after a month, a high percentage of new visitors in your analytics tool, add-to-carts rising even where purchases are not. An anchor ad at double the CPA target stays on as long as the campaign as a whole is holding target. Kill it and the efficient ads lose their volume.

Ads that convert and are still not ads. Andromeda will find the two people in the world who were going to buy from a weak creative, convert them, and show you a great CPA on almost no spend. A conversion with a bad CPC and no second signal is luck, not a performer. Check the add-to-cart rate and the click cost before you scale anything on a single conversion.

Fatigue versus Andromeda. Real fatigue has a trend you can watch: frequency climbs, CPM climbs, click-through rate falls over a week or two. An ad that worked for two days and died is not fatigue; the model exhausted the pocket it found. Treat the first as a creative refresh problem and the second as a concept problem, and do not report either as "the algorithm."

Retargeting after Advantage+

Turn Advantage+ audience expansion off for retargeting and tell Meta exactly who to reach. We trusted the algorithm with retargeting, and it answered by spending prospecting dollars on the people at the bottom of the funnel and calling it performance.

Three or four years ago retargeting meant handing Meta your pixel audiences and customer lists and saying go after these people. Under Advantage+ that list became a suggestion. Meta reaches them and then keeps going into new users, which is how frequency climbed on accounts that thought they were prospecting. The setting to

change is the one that reads use audience as a suggestion; switch it to further limit the audience.

What we run now:

- Manual retargeting ad sets aimed at add-to-carts from the last 180 days, plus any on-site action we have built audiences for, and Shopify Plus boosted audiences where the brand has them. Frequency is watched daily; it is the metric that tells you when this ad set is done.
- An existing-customer exclusion on every prospecting campaign, so the algorithm cannot fall back on your buyers when new customers get expensive. Your platform CPA will look worse. Your CAC will be real.
- Video-view audiences as a middle step. Collect the people who watched the first 25% of your best videos across several creatives into one audience, then retarget them with statics that state the value props plainly and carry an offer. It rebuilds the funnel the way we used to run it, on purpose, instead of leaving it to the model.

Advantage+ retargeting inside a CBO still has a place in some accounts, and it reaches a lot of new people. The question is whether those people convert at the efficiency a retargeting line is supposed to deliver. If you cannot answer that with new-visitor percentage and a CPMR trend, you are not retargeting, you are prospecting with a smaller audience and a bigger frequency.

It is always the offer

When the structure is right and the creative is fresh and the account still will not scale, it is the offer. We have watched a single ad in a single ad set run at \$2,000 for a month, and what it unlocked was not the ad.

Test the offer before you test the tenth creative. We ran 25% off against buy-two-get-one for a client this year and the BOGO roughly doubled revenue. BOGO used to lose for us. The macro reality is that every income bracket has less to spend, and people now want the most value out of the dollars they do spend. Offers that make the total feel bigger beat offers that make the price feel smaller.

Price tests under psychological barriers. Moving a product from \$200 to \$180 was the unlock for one brand. \$55 to \$49 does the same job. So does the opposite: \$6 a day out-performed \$5 a day for the same product, because \$5 reads as basic and \$6 reads as considered. Run the test on a landing page only, divert part of your Meta traffic to it, and judge it on new-customer CAC, not platform CPA.


Nobody sees your price change. Operators worry that raising a price from \$47.99 to \$49.99 will anger people who saw the old one. The user journey is too convoluted for that. Two percent of impressions click, those people land on a page, and only the ones who reach the PDP see the number. Run price and offer tests in a closed funnel from Meta only, exclude existing customers, and the brand risk is near zero.

The landing page is part of the offer. Changing the page is the cheapest performance refresh in the account. It changes who Meta reaches as well as who converts. The doctor-at-the-top example earlier is the extreme case; a quote and photo from the creator in the ad, or a page built around the one persona the ad speaks to, does a smaller version of the same thing. CPMs drop because the model sees unification between ad and page, and ads that have run for months get a second wind without a single new creative.

Positioning is the offer too. Two brands with the same product can sell at double the price when one has clean branding, one message, and a clear reason the price is what it is. Find three or four competitors for each element separately (branding, offer, messaging) and test the gap. Your ad brand guidelines do not have to be your website brand guidelines; a Meta ad and a billboard are different media and should look it.

The circle, and the 10-to-20% rule

The performance diagnostic circle



Change one lever. Read the funnel. Keep the next test written before launch.

When an account breaks you do not know which variable broke it, so you test the others. Structure, creative, copy, landing page, offer. If a change to one of them fixes performance, it was never the one you were sure it was.

We run that circle in order and we keep running it whether or not we think we found the answer, because the answer moves. The diagnostic that starts it is a funnel read, not a creative review: CPCs down but new-customer conversion rate down faster means an education problem on the page, not a creative problem. Platform CPA improving while CAC does not means new customers are not the ones converting, which calls for a structure and exclusion check before a creative one.

The rule. Spend 10 to 20% of every month's budget on things you have not tried: a campaign style, a persona, a page, an offer, a format. If the account is beating target you have the room. If the account is failing you have nothing to protect. The only accounts that should not be taking risk are the ones exactly on target, and there are no such accounts.

What risk is not. Leaving one creative live at a \$200 CPA and hoping is a risk, but a bad one. Risk is calculated: a \$100 to \$200 test of a campaign style you have wanted to try, a 10% budget shift to a new persona that becomes 15% if it works, a lo-fi batch you cut after one bad read and revisit in January when CPMs drop. Small pivots produce 10 to 20% improvements every six months. Big pivots, made on data, are the only way an account moves from losing to scaling.

Make it a system. Write down, before every test, what you will do if it wins and what you will do if it loses. Tag every ad by its attributes (format, hook, persona, background, person or no person, offer) and keep the data in one place so the next test is a hypothesis, not a guess. 90% of the brands we audit, at \$1M and at \$100M, cannot point to which hooks work in their own account. Organization is most of the edge.

The checklist

Run this against your own account before your next budget change. Every line is something we check in an audit.

- ☐ Prospecting campaigns exclude existing customers
- ☐ Retargeting ad sets have Advantage+ expansion set to further limit, with frequency watched daily
- ☐ CPMR and frequency are on the daily dashboard, next to CPA and CAC
- ☐ Ads are judged on a 2-to-3-day window with a written pause rule tied to how much room the account has
- ☐ Anchor ads are identified by CPC, new-visitor share and frequency, and protected from CPA-based pausing
- ☐ Every new ad set carries 8 to 10 ads across formats so entity IDs do not collapse it into one ad
- ☐ Campaigns are split by margin profile or product line, and the split is re-tested when growth slows
- ☐ 10 to 20% of this month's budget is committed to something untested
- ☐ Pixel and Conversions API fire once per event (we find double and triple firing in \$20M accounts)
- ☐ One offer or price test is running on a closed landing page right now

What happens next

If more than three of those boxes are unchecked, the fastest fix is to let us look. DigiCom runs a free audit for brands spending \$25,000 or more a month: 10 to 12 hours of a senior strategist on your account covering structure, settings, pixel setup, creative tagging, landing pages and where every dollar actually lands. It ends in one of three honest answers: yes we can hit your goal, yes with specific changes, or no. We cap the number we run each month, and the team that runs the audit is the team that runs the account.

Take the two-minute Growth Leak Check at digicom.io/audit. It tells you which of the four Andromeda problems you have and whether you qualify for an audit slot this month.

THE PEOPLE BEHIND THE PLAYBOOK

Meet DigiCom.

The no bullshit growth agency.

Born in New York, powered by global talent. We help ambitious DTC brands build growth engines through paid media, creative, CRO, influencer partnerships and measurement. We care about profitable growth, not vanity metrics.

\$200M+

MEDIA SPEND

\$1B+

REVENUE GENERATED

200+

BRANDS

20+

PLATFORMS

2X INC. 5000 HONOREE**2025 / 2026**

Google Partner | Meta Business Partner | TikTok Marketing Partner

MORE FROM THE TEAM

[Explore 60+ case studies >](#)[Meet the DigiCom team >](#)[Explore digicom.io >](#)

Metrics reflect DigiCom's published agency / leadership track record.

NEXT STEP

Find the growth leak.

DigiCom's free no-BS audit looks at the things that quietly stop good brands from scaling: structure, measurement, creative and landing pages.

REQUEST YOUR FREE AUDIT >

Senior-team review. Straight answers.



SCAN TO VISIT

For qualifying brands spending \$25K+ per month, as described in this paper.

KEEP THE CONVERSATION GOING

Performance Anxiety + DigiCom growth conversations

Real operators. Actual numbers. Useful experiments.

[Watch the podcast playlist >](#)



DigiCom | October 2026 | digicom.io