

CREATIVE / META ADS

The Andromeda creative playbook

The research, hooks and testing system behind direct-response creative.

Turn customer language, persona diversity and performance data into the next winning ad.

THE PLAYBOOK IN ONE LINE

RESEARCH. DIVERSIFY. TEST. LEARN.

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

The executive read.

What matters, before you touch a budget.

01

Creative is the new targeting

Variations in persona and awareness matter more than cosmetic swaps.

02

Write in customer language

Use reviews, comments, Reddit and survey objections as your raw material.

03

Solve for the first seconds

Test visual, audio and text hooks; diagnose the drop-off by metric.

04

Create an ad bible

Tag every asset and plan the next test before the current one ends.

THE FIELD GUIDE

01 Creative is the targeting now

02 Sell the problem, not the product

03 Research: where the language comes from

04 The persona by awareness matrix

05 Hooks first

06 Formats and where the content comes from

07 The ad bible

08 Using AI for the research and the critique

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.

Creative is the targeting now

Media buying is not what it was. Advantage+, Andromeda and every platform's move toward automated delivery mean the ad set no longer decides who sees your ad; the creative does. Pre-2020 you could run the same creatives for months and print money. Now you have one to three seconds to stop a scroll, the model groups look-alike ads into one, and a winner that carried an account for weeks can stop overnight.

This paper is the direct-response creative process DigiCom runs for the brands it manages: how we find the problem the product solves and the person who has it, how we research the language they use, how we build creative diversity that goes beyond swapping a headline, how we test hooks before anything else, and how we tag every ad so the next test is a hypothesis instead of a guess. The last section is how we use AI for the research and the critique, which is where it is actually useful.

It is written for operators spending \$25,000 or more a month who are doing creative volume and still not seeing it move the account. Volume without diversity hits the same people. Diversity without research says nothing to anyone. The playbook is both.

By Hemant Varshney, Danielle Dror and Mike Constandatos, co-founders of DigiCom, with Jake Orenstein, Director of Growth.

Sell the problem, not the product

You are not selling a product. You are selling a problem, to the specific person who has it, in the state of mind they are in when the ad reaches them. Everything else in this playbook follows from getting those three things right.

Direct response is stimulus. You are trying to make someone stop what they are doing and act in that moment. That is a psychological and behavioral event, not a design one. The first five seconds earn the stop; the rest of the ad has to keep stimulating through to the click. A polished, on-brand creative that does not do that is not a brand asset, it is an expensive way to be ignored.

The persona is a hypothesis. Brands come to us with a target customer. We find out, from the data, that it is someone else, often a different gender, age or life situation, and sometimes two or three different people depending on the product. A CPG food brand wanted health-conscious moms of two; the persona that also converted was the gym-goer who wanted to eat it without breaking their diet. Both worked, from different creatives. Write the persona down as what you believe, then let purchases rewrite it, and expect it to change again with every product launch.

How niche to go depends on where the product sits. We use Maslow's hierarchy to decide. A product that meets a physiological or safety need (water, basic supplements, protection) can be sold on demographics; everyone needs it, so the persona is broad and the message is the need. The further up the pyramid the product sits, toward belonging, esteem and self-actualization, the more the persona has to be a subculture and not a demographic: not every woman over 30, but the specific person who wants to belong to the group this brand signals. A novelty kids' soap is not sold to parents; it is sold to the parent who wants bath time to stop being a fight.

Nobody buys ingredients. Features are what the product has. The ad is about what the person feels before, and what they will feel after. Someone having a bad day is looking for the thing that fixes the day; someone excited buys on impulse. You cannot know which one you have, so the ad shows the after-state and lets the person who

needs it recognize themselves.

Research: where the language comes from

The words in a winning ad are never ours. They come from customers, and the research job is to find the specific sentence someone has already said about the problem and put it in the hook.


Source	What you are looking for	What to ignore
Your reviews	The rare review that names a specific situation ("finally sleeps through the night"), not the outcome everyone names	"Great product, life-changing"
Comments on your best ads	Skepticism and objections. If 15 to 20% of commenters say one thing, 20% of your audience is thinking it	Praise, tags, spam
Reddit	Threads about the problem, not the product. It is anonymous, so people say what they mean	Nothing: it does not have to be about your product
Post-purchase surveys	Why now, what else they tried, who they bought it for	Which channel they "heard about you" on
TikTok search on the problem	The slang and phrasing people actually use when they talk about it	Views and likes
Ads Library	Concepts and angles that are trending across your category; nobody is reinventing the wheel	Copying the ad itself

Answer the objection in the ad. For a \$599 kitchen appliance, Reddit told us the price was the first thing 30 to 40% of people hit. So the ads said it: it is not cheap, and here is everything it does for that money. Addressing the objection head-on converted the people who were already thinking it, and it came from a thread we did not write.

Each platform has a voice. Reddit speaks in TL;DR and AITA; TikTok in trends and sounds; Meta in talking heads and product-in-hand; YouTube Demand Gen wants the product laid out and the graphic shown, not a creator talking, which we learned when an apparel brand's winning Meta creatives flopped there. Before the persona and the subculture, decide where the ad is being served and write for that room.

Comments are a microcosm. Ads with real spend and longevity behind them accumulate comments, and those comments are the cheapest focus group you will ever run. Read them on your top ads regularly. The objection that keeps showing up is your next hook.

The persona by awareness matrix



Creative diversity is not six formats of the same message. It is a grid: personas across the top, awareness stages down the side, and a different ad in every cell.

If all your creative is solution-aware ("here is the product that fixes X") you are saying the same thing to the same people with different creators, and your frequency climbs while your reach stalls. Andromeda makes this worse because it groups those look-alike ads into one entity. The way out is up the funnel: creative for people who know they have the problem but not that a solution exists, and creative for people who have not named the problem yet.

	Problem-unaware	Problem-aware	Solution-aware	Product-aware
The health nut (gym at 6am, reads labels)	What your morning routine is missing	Why your post-workout energy crashes	The ingredient that fixes it, compared	Why ours over the one you have
The low-tox mom (switching the household)	Old me / new me: the drugstore version next to ours	What is in the one you are using	Our product on the vanity, nothing else	Reviews from moms, the bundle offer
The skeptic (saw the ad, did not click)	—	The Reddit objection, answered	Ingredients panel as the hero image	Price justified line by line

Two things we learned building these grids across accounts:

- The cells have different metrics and that is correct. One creative sits at an 8% conversion rate and a 0.55% click-through rate; another at a 2.5% click-through rate with the best CPCs in the account and one of the lowest conversion rates. Neither gets turned off. They are doing different jobs in different cells, and the account only scales because both exist.
- Top-of-funnel cells need to be understood in one frame. The ad that scaled fastest to new audiences for one brand was an old-me / new-me static: a drugstore product next to ours, and a line about which hotels stock it.

People understood it instantly, so it scaled at the lowest cost per click in the account. The value-prop ads (what is in it, what is not) needed 2 to 2.5 exposures to land and belonged in the middle of the funnel. Same product, same brand, different cells, different frequency budgets.

Build the grid before the brief. If a concept does not sit in a cell, it is a format variation, not diversity, and it will be grouped with its siblings the moment it uploads.

Hooks first

Use the drop-off to find the problem.

01

HOOK RATE

Did they stop?

02

WATCH TIME

Did they stay?

03

CTR

Did they click?

04

CVR

Did they buy?

A strong metric upstream and a weak one downstream points to the next thing to fix.

If the first one to three seconds lose them, nothing after it matters. Most of our testing budget goes to the hook, because it moves the needle more than anything else and it is the cheapest thing to vary.

Three hooks in every video. The visual hook is what the first frames show: a face, the product, an unboxing, a before. The audio hook is the first sentence spoken. The text hook is the overlay, and it is how you make one video speak to three pain points without reshooting. Test them separately; a video with a dead visual hook and a great audio hook reads as a dead video until you split them.

Same message, different sentence. A kids' soap brand ran one concept, making bath time fun, across a batch of ads. One spent tens of thousands of dollars; the rest flopped. The message was identical. The difference was "mom hack: make bath time fun with one product" against "POV: bath time is now your kid's favorite part of the day." When a message fails, test the phrasing before you abandon the message.


Hook rate before everything. If we could track one creative metric it would be hook rate (three-second views over impressions). Read it first, then average watch time, then click-through, then conversion, and let the drop-off tell you which part of the ad is broken: a high hook rate with low watch time is a body problem, a high click-through with a bad conversion rate is an ad that promises something the page does not deliver. On YouTube Demand Gen the equivalent read is the first 25% of the video.

Stretch the asset. A creative that did not work is often a creative that was cut wrong. Before you shelve it, re-hook it, shorten it, add text, change the first frame. Most content has more longevity than the first test shows, and a brand with limited video cannot afford to throw a shoot away on one hook.

Formats and where the content comes from

Same video. Different handle. Different hook rate.

CREATOR HANDLE



BRAND HANDLE



Reported hook rate from one same-video / same-broad-audience test in the manuscript.

The best-performing ad in one new account was shot sideways on an iPad, with sauce on the stove and a wooden spoon in frame. The brand's contact called it lame. It spent for months and became 17 iterations. Polish is not the variable.

UGC versus polished. Brands kick off wanting the on-brand video. We ask for \$1,000 to test the ugly one. Most of the time the ugly one wins, because it looks like the feed instead of an ad, and the brand moves its content budget toward founder and creator video for the rest of the engagement. Video is still the format that feels most native and carries the story; it is the hardest thing to run an account without.

Influencers after Andromeda. Two years ago we told brands to buy volume UGC over one influencer. That is no longer true for every brand. Partnership ads run from a real creator's handle are finding audiences we cannot reach from the brand handle: the same hook from the same video got a 30% hook rate from the influencer's handle and 13% from the brand's, in the same broad audience. The theory is that the model uses the handle to find like-minded people, not the followers themselves. Mass UGC platforms, by contrast, give us a usable creative under 5% of the time; the content reads as people after free product. Build a small bench of creators who actually use the product. That relationship is the difference between a creative that works and a pile that do not.

Statics are back, carousels are underrated. For a brand with no video, statics can carry the top of the funnel if the frame makes the product clear in one look: the comparison, the product on a plain background or a vanity rather than a beach, old-me / new-me. Carousels tell a story the way video does and are the best substitute when you cannot film. Analysis across one account showed flat-lay and vanity backgrounds scaling at lower frequency than lifestyle settings; that kind of finding only shows up if the attributes are tagged.

Lo-fi and AI. Lo-fi and AI-generated video are getting good enough to pass, and a text-over-waves ad in the Calm style costs nothing to make. They are also a Q4 risk: one lo-fi batch ran at \$70 against a \$40 target and was shelved until January, when CPMs drop. Use AI to get a concept out of your head and in front of the design team as an image, and to generate the formats you are missing, not as the finished ad.

The ad bible

Tag every attribute of every ad, join it to the performance data, and pivot. It is the most tedious thing we do and it is where the playbook for the next quarter comes from.

What gets tagged. For each ad, the link and then ten or so attributes: format (video, static, carousel, GIF), length, hook type (visual, audio, text), persona, awareness stage, person in frame or not, product position, background, setting, offer, creator versus brand handle, and the trend it borrowed (us-versus-them, POV, comparison). Then pivot against click-through, hook rate, CPA, CAC and frequency. The findings read like this: a human holding the product beats no human; blue backgrounds beat white; product left and person right converts better. Whatever yours turn out to be, that becomes the brief for the next batch, and it will be wrong again in two months, which is why the sheet is never finished.

Reading a batch. Never test one or two ads; you learn nothing. A batch is many variations in a clean ad set. For video, read hook rate, average watch time, click-through and conversion rate against each other. High hook rate and low watch: the body is wrong, keep the opener. High click-through and low conversion: the ad promises what the page does not deliver, fix the match. Product visual versus person in the first frame: compare the two hook rates and eliminate the loser from the next brief.

The roadmap, written before the test. Every test has two next steps written down in advance: if it wins, here is what we scale and iterate; if it loses, here is what we learn and what we try instead. The alternative is throwing concepts at the wall, one sticks, it fatigues, and you are playing catch-up with nothing behind it. Pull on the thread that is working and close doors behind you: if orange outperforms black, you stop making black, your floor rises, and the next test starts from a higher base.

Organization is the edge. Across brands at \$1M and \$100M, more than 90% cannot say which hooks work in their account. Video assets in one folder, statics in another, tagged by persona and hook, with the data joined, is the whole system. Our hit rate on new creative is not 50%, but it climbs the longer we run the bible on an account, and a move from 5% to 10% is a doubling of every creative dollar.

Using AI for the research and the critique

AI writes bad ad copy and excellent research. The brands getting value from it are the ones feeding it their own data and asking it why, not the ones asking it for hooks.

One GPT per product, fed with performance. Build a custom GPT (or Claude project) per product, not per brand, and load it with the brand guidelines, the assets, the reviews and surveys, and the actual numbers: the best copy by click-through on each platform and the worst copy too, the ad bible with its tags, the charts. It is primed on what worked in your account instead of what worked on the internet. Give it a role and a task: analyze these two pieces of copy as an Ivy League marketing professor and tell me the five variables that make the first one work. The answer, in-group versus out-group framing, belonging, specificity, is a list of variables you can now test on purpose.

Business analyst, not growth marketer. Our internal engine took 49 versions of the prompt to get useful. The change that mattered was the role: asked to act as a growth marketer it returned five-word bullets; asked to act as

a business analyst with a framework for causation versus correlation (spend rose on this campaign, conversions rose, is conversion rate the cause) it returned analysis we could act on. The framework is the prompt.

Critique, not generation. The most useful question is the negative one. I hate this creative, why? It comes back with the color clash, the off-center composition, the words that are not drawn out, and the design team gets a brief instead of a feeling. The same works for repositioning: here is the brand, here is who they want, here is why the current branding reads cheaper than that, build the brand book within their colors and fonts. That is the opposite of asking it for the destination, and it is where it is good.

Expect 10 to 30 prompts. You are the editor. Whatever it produces gets rewritten in the language from the research section, in the register of the platform, by a person who knows how the customer talks. The flowery copy it defaults to is not direct response; this is not Mad Men, and the ad that wins is usually the plainest sentence a real customer already said.

The checklist

- ☐ Each product has a written persona with the evidence it came from, and a date it was last revised against purchase data
- ☐ A persona by awareness grid exists and every live concept sits in a cell
- ☐ Top five ads' comments are read weekly and the top objection has an ad answering it
- ☐ Every video is tested with more than one hook (visual, audio or text) before the concept is judged
- ☐ Hook rate, watch time, click-through and conversion rate are read against each other per batch
- ☐ Ad sets launch with 8 to 10 ads across formats, not variations of one
- ☐ One founder or creator video shot on a phone is live in the account right now
- ☐ Every ad is tagged with its attributes and joined to performance data in one sheet
- ☐ Every test has its win and lose next step written before launch
- ☐ A per-product GPT holds the brand assets and the real performance data

What happens next

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, 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 of audits each month, and the team that runs the audit builds the creative, writes the copy and buys the media.

Take the two-minute Growth Leak Check at digicom.io/audit to see whether you qualify for a slot this month.

A free audit. Three possible answers.

YES

We can hit your goal.

YES, IF...

We fix specific leaks first.

NO

We will tell you why.

[See how the senior-team review works >](#)

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