

Meta Advertising Audit and 90-Day Scale Plan

Sample deliverable — client identifiers redacted

Client: [REDACTED] — DTC pet skincare brand, topical range

Account reviewed: [REDACTED]

Media spend reviewed: approximately \$145,000 across the review window

Review window: 1 January to 30 September 2026

Prepared by: Serafinity

Spend rose four times between the first quarter and the third. Revenue rose twice. Every extra dollar returned \$0.70, and the ads were never the reason.

About this sample. This is a real audit delivered to a real client, published with their written permission and anonymised at their request. Brand name, domain, account and dataset identifiers, campaign, creator and asset names, product names and competitor names have been redacted. Currency totals are rounded. Every percentage, ratio, benchmark and conversion rate is exactly as delivered. Nothing in the analysis has been softened, removed or improved after the fact.

Contents

1. Scope and method

What was examined, and over what period

2. Primary finding: spend rose 4x and profitability fell

The single number that explains the year

3. Opportunity Score and account health

What Meta recommends, and what it misses

4. The campaign timeline: four structures in nine months

What ran, when, and what each returned

5. Funnel conversion rates, stage by stage

Every step, and which steps can be trusted

6. Measurement and attribution

Three attribution windows and three events that stopped adding up

7. Unit economics

Order value against acquisition cost

8. Retargeting and retention

What past visitors and customers are being shown

9. Creative inventory and performance

Six findings, including two switched-off winners

10. Competitor activity in the US market

Who is buying this auction, and how

11. Unclaimed positioning

Five positions available now

12. 90-day remediation and scale plan

Week by week

13. Budget scaling thresholds

\$18k to \$100k, gated

14. Creative concept territories

Seven territories, each tied to a finding

15. Creative production model

Who makes what

16. Seasonal planning

Peak season and allergy season

17. Projected outcomes, timelines and risks

Including what looks worse first

18. Reporting and working cadence

What arrives, and when

19. Summary of findings

The whole audit on one page

1. Scope and method

This audit is based on primary data pulled directly from the client's Meta advertising account and its pixel dataset. Every ROAS, revenue and purchase figure is Meta-reported and has not been reconciled against the store's own order data, because store access was outside the scope of this review.

What was examined

- All 18 campaigns that delivered spend in the ad account between 1 January and 30 September 2026.
- Campaign, ad set and ad-level delivery and conversion metrics: spend, impressions, reach, frequency, CTR, CPM, ROAS, cost per purchase, and every funnel event from link click through purchase, by month and by week.
- Ad set configuration: optimisation goal, attribution setting, budget type and bid strategy for every ad set with spend.
- Video engagement at ad level, plays and 25% retention, for the 28 highest-spend ads.
- The pixel dataset: event match quality and match-key coverage across 13 events.
- Meta's Opportunity Score and its six live recommendations.
- The live Meta Ad Library for the United States on 2 October 2026, active ads only, across several competitors.
- The storefront: product range, price points, bundles, thresholds and subscription offer.

What was not examined

Store orders, the subscription platform, email and SMS flows, and margin data.

2. Primary finding: spend rose 4x, and each added dollar returned \$0.70

In the first quarter the account spent \$20,500 and Meta attributed \$42,900 of revenue to it. In the third quarter it spent \$82,100 and Meta attributed \$86,100.

Quarter	Spend	Revenue	ROAS	Purchases	Cost per purchase	CPM	Cost per 1,000 reached
Q1 (Jan-Mar)	\$20,500	\$42,900	2.10	836	\$24.48	\$27.25	\$54.83
Q2 (Apr-Jun)	\$42,400	\$59,800	1.41	1,194	\$35.48	\$46.63	\$113.71
Q3 (Jul-Sep)	\$82,100	\$86,100	1.05	1,800	\$45.63	\$42.91	\$142.45
Nine months	\$145,000	\$188,800	1.30	3,830	\$37.85	\$40.56	\$109.63

Month by month

Month	Spend	Revenue	ROAS	Cost per purchase	CPM	Frequency
January	\$10,400	\$21,300	2.06	\$25.05	\$24.76	2.04
February	\$4,600	\$9,600	2.09	\$23.25	\$26.18	2.03
March	\$5,500	\$12,000	2.18	\$24.51	\$35.15	1.93
April	\$9,500	\$11,500	1.22	\$46.27	\$45.94	2.16
May	\$15,700	\$23,800	1.51	\$32.87	\$47.86	2.63
June	\$17,100	\$24,500	1.43	\$33.58	\$45.93	2.45
July	\$22,600	\$28,000	1.24	\$40.05	\$46.05	3.13
August	\$28,600	\$25,200	0.88	\$53.16	\$43.15	3.52
September	\$30,900	\$32,900	1.07	\$44.33	\$40.66	3.29

Why this matters

ROAS halved between Q1 and Q3. Breaking that fall into its parts shows where it came from.

- CPM rose 57%, from \$27.25 to \$42.91. This is the largest single driver by a wide margin.
- Click-to-purchase rate fell 12%, from 6.50% to 5.69% of link clicks.
- Average order value fell 7%, from \$51.32 to \$47.82.
- Link click-through rate fell 3%, from 1.71% to 1.65%.

The ads did not stop working. People clicked at almost the same rate in September as in January. What changed is the price paid to put the ads in front of them, and nothing downstream improved to pay for it.

Reach and frequency explain the rest. Between March and September, monthly spend rose 5.6x, from \$5,500 to \$30,900. Frequency went from 1.93 to 3.29. The cost of reaching 1,000 people went from \$54.83 in Q1 to \$142.45 in Q3. The extra budget was not spent finding new customers. Most of it was spent showing more ads, at a higher price, to a pool that grew far more slowly than the budget did.

3. Opportunity Score and account health

Meta's Opportunity Score for the account is 84 out of 100. That is a healthy score. Accounts with structural problems typically sit in the 50s and 60s with a long list of fixes. This account generated six recommendations, and none of them touch the reason ROAS halved.

Metric (1 Jan to 30 Sep 2026)	Value
Total spend	\$145,000
Meta-attributed revenue	\$188,800
Meta-reported blended ROAS	1.30
Purchases	3,830
Cost per purchase	\$37.85
Average order value	\$49.29
September spend	\$30,900
Opportunity Score	84 / 100

A note on budget

Active campaigns hold large budgets, and if those budgets fill, spend rises further. Two of the three use cost-cap and minimum-ROAS bidding, which currently hold delivery below budget. Either way, the account is configured to spend more, not less, than the month that returned 1.07.

What Meta is actually recommending

- **Send more match keys for Purchase** (4 points, 16 ad sets). Meta wants a login ID passed with purchase events. This connects to section 6: the click ID is present on only 44.4% of purchase events. Worth investigating.
- **Turn on one more Advantage+ creative enhancement** (4 points). Meta estimates 6% lower cost per result from text variations. Low risk and worth testing, though we have not seen a significant difference in performance elsewhere.
- **Switch seven ad sets to Maximise Value of Conversions** (1 point each, two recommendations). Meta estimates 7% higher ROAS. The account already has evidence here: the largest campaign of the year ran value optimisation for \$74,600 and returned 1.08. On a catalogue where most products sit between \$23 and \$39, there is little value spread for the system to chase. It should still be tested, because we have seen large differences between optimising for value and for purchase count.
- **Advantage+ audience and mixed formats on the engagement ad set** (1 point each). We would not act on either. Section 4 recommends switching that campaign off.

The score on its own says little. Read alongside the frequency problem and Meta's own recommendation to change delivery, it is clear the current setup is not working. The account needs audience exclusions, a real search for new customers, and a return to automatic delivery with value settings tested properly.

4. The campaign timeline: four structures in nine months

Campaign names in the account fall into four distinct naming systems, each with its own structure, optimisation goal and attribution window. Grouping spend by system gives the clearest single view of the year.

Structure	Active	Spend	ROAS	Cost per purchase	Optimised for	Attribution
A. Original scaling CBO plus ad-set-budget testing	Jan-Jul	\$38,900	1.82	\$28.16	Purchase count	7-day click
B. Creator testing, ABO plus four product tests	Mar-Jul	\$18,200	1.31	\$36.85	Purchase count	7-day click
C. Value-optimised prospecting CBOs, promotional sale, retargeting	Jun-Sep	\$84,400	1.09	\$44.25	Purchase value	7-day click plus 1-day view
D. Cost cap, two Advantage+ shopping campaigns, page engagement	Sep onward	\$3,200	1.14 on sales campaigns	\$42.32	Purchase count	7-day click

Two small engagement and boosted-post campaigns (\$175) sit outside these groups. Structure D's ROAS excludes its \$1,134 engagement campaign, which records no purchases.

Observations

- **The best-performing structure of the year was switched off.** Structure A returned 1.82 on \$38,900 at a \$28.16 cost per purchase. Its main ad set alone took \$29,500 at 1.77 with a \$31.21 CPM. It spent \$366 in July and nothing after.
- **Its replacement cost 57% more per purchase.** Structure C took \$84,400, which is 58% of the year's spend, and returned 1.09 at \$44.25 per purchase.
- **The comparison flatters the replacement.** Structure A ran on 7-day click. Structure C ran on 7-day click plus 1-day view, which also counts people who saw an ad, did not click, and bought within a day.
- **April is where the cost base moved.** CPM went from \$35.15 in March to \$45.94 in April and never came back. Structure B opened that month at a \$74.16 CPM on small daily budgets across creator ad sets.
- **Every handover reset the learning.** Each new structure launched new campaigns, new ad sets and new ad IDs for creative that already had history. The same before-and-after video appears under three ad IDs in three campaigns.
- **A fourth rebuild began three weeks ago.** Structure D went live on 8 and 30 September on a third bidding approach. Early numbers are too thin to judge and are not a signal in either direction.
- **On value optimisation:** the value-optimised structure is the worst-performing one, but the likely reason is that it was not set up correctly rather than that value optimisation is wrong for this account.

5. Funnel conversion rates, stage by stage

The table traces every recorded event across the nine months. The verdict column matters more than usual here, because several steps cannot be read at face value.

Stage	Volume	Step rate	Typical range	Verdict
Link clicks	59,042			
Landing page views	46,792	79.3%	70–85%	Normal
View content	60,361	129% of landing page views	n/a	Multiple views per visit
Add to cart	8,825	14.6% of content views	10–20%	Normal
Initiated checkout	2,694	30.5%	50–65%	Not reliable
Purchase	3,830	142% of checkouts	60–75%	Not possible

Reading the table

The account records 3,830 purchases and 2,694 initiated checkouts. More people cannot finish a checkout than start one. This is not a rounding issue. Purchases have exceeded checkouts in every month since March, and in August and September the ratio reached 1.9 purchases for every checkout recorded. That means the checkout step is being under-counted, the purchase step is being over-counted, or both. Until that is resolved, the two rates most audits lean on, cart-to-checkout and checkout-to-purchase, cannot be used to diagnose this store.

The steps that can be trusted

Rate	Q1	Q2	Q3
Link click to purchase	6.50%	8.23%	5.69%
Add to cart to checkout (as recorded)	41.7%	45.0%	21.0%
Purchases per checkout (as recorded)	1.09	1.39	1.69
Average order value	\$51.32	\$50.09	\$47.82

- Link click to purchase is the cleanest read on the store. It ran at 6.5% to 8.2% in the first half and fell to 5.69% in Q3. In August it was 4.78%. A click-to-purchase rate above 5% is strong for cold traffic, which is itself a hint that a share of the traffic is not cold. Section 6 picks this up.
- Recorded cart-to-checkout halved in August. It held between 39% and 48% from January to June, then fell to 17.7% in August and 16.4% in September. Add-to-cart volume nearly doubled over the same weeks while checkouts stayed flat. A real behavioural change that sharp is unlikely. A change in what fires the add-to-cart event is more likely.
- January and February landing page views are under-recorded. Only 42% and 57% of link clicks registered a landing page view in those months, against 82% to 91% from March onward. Something in the pixel setup changed in March.

6. Measurement and attribution

Nothing in this section says the tracking is broken beyond use. The purchase event itself is in good shape on paper. What it says is that the numbers used to judge the account this year were measured three different ways, and that three events stopped adding up at specific points.

What is working

Event	Match quality (of 10)	Click ID coverage	Read
Purchase	9.3	44.4%	Strong identity match, weak click match
Initiate checkout	8.3	50.0%	Good
Add to cart	7.2	67.7%	Acceptable
View content	7.1	48.4%	Acceptable
Page view	7.0	50.2%	Acceptable

A 9.3 on Purchase means email, phone, name and address arrive with every order. That is as good as it gets. The weak point is the click ID, present on fewer than half of purchase events. Without it, Meta leans on identity matching to connect an order to an ad, which is more generous to view-through and to existing customers.

Finding 1: three attribution windows in one year

- Structure A ad sets: 7-day click, with 1-day view explicitly removed.
- Structure C ad sets: 7-day click plus 1-day view, optimising for purchase value.
- Structure D ad sets: back to 7-day click.

The account-level ROAS for each month is a blend of whichever windows were live. The 2.10 in Q1 was earned on click-only attribution. The 1.05 in Q3 was earned with view-through credit included. The real decline is larger than the reported one.

Finding 2: new and existing customers are reported as one number

The pixel already receives three separate events: purchase, new-customer purchase and existing-customer purchase. It also receives a subscription event. So the store can already tell Meta which orders are first orders. No campaign in the account before 8 September carries a customer exclusion, and every ROAS figure for the year counts repeat customers alongside new ones.

For a consumable product with a subscription offer, this matters a great deal. A 3.3 frequency, a click-to-purchase rate above 5% and a value-optimised campaign with no exclusions are what it looks like when prospecting budget drifts toward people who already buy.

Finding 3: events that stopped adding up

When	What changed	Most likely cause
March	Landing page views jump from 42-57% of clicks to 88%	Pixel or consent configuration change
March onward	Purchases exceed initiated checkouts every month	Checkout event firing in the browser only while purchases arrive server-side, or express checkout skipping the event
Late August	Add-to-cart volume doubles, checkouts flat	A cart or bundle element firing add-to-cart on each click. The cart drawer's bundle-builder buttons are the first place to look
Late September	Content views fall from 1.2-1.4 per landing page view to 1.05, then to 0.39 on 1 October	Traffic sent to a page that does not fire the event, or a theme change

Causes are hypotheses drawn from the pattern of the data. Each is testable in under an hour with Events Manager and a test order.

Why this matters for delivery

Meta's system optimises on the events it receives. If add-to-cart is inflated and checkout is under-counted, the mid-funnel signal it learns from is noise. If repeat customers count the same as new ones, it will find repeat customers, because they are the cheapest purchases available. That changes who the ads are shown to, which is why this is the first thing to investigate.

7. Unit economics

The average Meta-attributed order over the window was \$49.29. In Q3 it was \$47.82, and each one cost \$45.63 to acquire. That leaves \$2.19 to cover the product, the packaging, shipping, payment fees and everything else.

Product	Price
Hero spray	From \$23, \$28 in the bundle builder
Rinse	\$26
Shampoo	\$26
Travel spray	\$24
Wipes	\$14 and \$39
Two-product set	From \$51
Free shipping threshold	\$60
Free gift threshold	\$80
Bundle builder	Choose 3, save 10%

Two consequences

First, a single-product first order cannot carry a \$45 acquisition cost. The hero product sells for \$23 to \$28. On first-order economics alone, the account stopped being viable somewhere between March and April.

Second, the business case rests on what happens after the first order. The site sells subscriptions, runs a loyalty programme, and the product is a consumable that runs out. That is the right model for this category. It also means the number that decides how much the brand can pay for a customer is 90-day customer value, and that number does not live in the ad account. It is the first figure we need, alongside contribution margin and the split of new against returning orders.

Recommendation

Send paid traffic to the two-product set and the bundles rather than to single bottles. Make the distance to free shipping visible on the product page and in the cart. Test a first-order bundle priced just above \$60. Put the subscription option in front of first-time buyers rather than after the fact.

8. Retargeting and retention

Across nine months and \$145,000 of spend, the account spent \$194.68 on retargeting. That is 0.13% of spend.

Metric	Retargeting ad set	Read
Spend	\$194.68	July and August only
ROAS	1.48	Four purchases only
Reach	529	Very small pool
Frequency	6.33	8.50 on the main ad set
Audience	Cart and checkout, 180 days	One segment, one window

Why this did not work

A 180-day cart audience that reaches 529 people was built wrong, or had its exclusions set too tightly. The account recorded 8,825 add-to-cart events in the window. Even after discounting for the inflation described in section 6, several thousand people put a product in a cart this year and did not buy. They have received almost nothing.

In fairness, broad Advantage+ campaigns do some retargeting on their own, and part of the 3.3 frequency in Q3 is exactly that. But it is uncontrolled: the same prospecting ad is shown again to someone who already knows the product, at prospecting CPMs, with no change of message.

For this brand, warm audiences should be the cheapest revenue in the account, for three reasons specific to the product.

- The purchase is a considered one for a worried owner. People compare against the vet visit and the prescription before they buy. A second and third touch with proof does real work.
- The product runs out. A bottle used daily on a recurring problem has a predictable refill window.
- The range is a routine. A buyer of the hero spray is the natural next customer for the rinse and the wipes.

9. Creative inventory and performance

The table covers the highest-spend ads of the year. It shows where the account's proof actually sits.

Ad	Spend	ROAS	Cost per purchase	CTR	Hook
A — creator video, science angle	\$18,800	1.35	\$32.24	2.73%	10.1%
B — vet point-of-view video	\$8,400	1.96	\$23.57	2.08%	63.5%
C — in-house before and after (three ad IDs)	\$8,300	2.11	\$31.27	3.15–3.84%	66–73%
D — creator video, rinse (Apr)	\$5,800	1.92	\$24.27	1.79%	30.6%
E — comparison static, active ingredient against alternative	\$4,400	1.15	\$42.66	1.99%	n/a
F — creator video, rinse, product-page cut	\$4,100	0.84	\$53.51	1.97%	12.4%
G — testimonial with before and after, new hook	\$3,800	1.59	\$29.10	2.18%	6.0%
H — creator video, two products (May)	\$3,700	1.45	\$40.67	1.61%	11.4%
I — creator video, science angle, August re-cut	\$3,600	0.76	\$56.18	2.22%	8.8%
J — creator video, hero spray, product-page cut	\$3,500	1.21	\$35.86	1.46%	6.5%
K — before and after, second copy	\$3,200	0.93	\$46.36	2.47%	9.2%
L — creator collaboration video (Aug)	\$2,600	0.52	\$81.15	2.46%	6.2%

Hook = 25% video retention as a share of video plays. On short clips this figure runs high, so compare ads of similar length rather than reading it as an absolute.

Finding 1: the two most efficient assets of the year are switched off

Asset C returned 2.11 on \$8,300 across three ad IDs, with the highest click-through rate in the account. Asset B returned 1.96 on \$8,400 at a \$23.57 cost per purchase. Both sat in Structure A and stopped spending when those campaigns were paused. Relaunching them in the new structure is the first creative action.

Finding 2: concentration risk

One ad took 13.0% of the year's spend. One creator, across five ad IDs, took 19.7%. The lead ad returned 1.35, which is respectable, and its two August follow-ups returned 0.76 and 0.67. When the next version of a winner returns half as much, the account does not have assets that can scale. It has one video that worked, and a lot of risk if it stops.

Finding 3: the same asset performs differently by structure

Asset C returned 2.45 in Structure A, 2.39 in the testing campaign, and 1.65 as a copy inside Structure C. Same video, same product, same landing page. What changed was the campaign around it: the optimisation goal, the attribution window, and a budget shared with twenty creator ad sets. This is the strongest evidence in the account that structure, not creative quality, drove the Q3 result.

Finding 4: creator content opens weakly

Eight of the twelve ads above hold fewer than 13% of viewers to the quarter mark. The demonstration-led assets hold 60% or more. Even allowing for length, the gap is wide. Creator videos in this account tend to open on a person talking. The assets that work open on the problem itself.

Finding 5: almost no static creative

Two statics appear in the top 28 ads by spend. The larger is the active-ingredient comparison, at 1.15 on \$4,400. It is the only ad in the account that explains why the product is different in a single frame. The format is under-tested. Meta's own recommendation engine flags the same gap. Static, video and carousel formats all need testing.

A note on claims and platform policy

Several ad names and testimonials reference infection, healing and treatment. Before-and-after imagery and treatment language for an animal topical carry policy and regulatory risk. Nothing in the account indicates enforcement to date. We would move the copy toward relief, soothing and skin-barrier support, which is the language already used on the storefront, and keep the customer photographs.

10. Competitor activity in the US market

A live sweep of the Meta Ad Library on 2 October 2026, filtered to active ads in the United States.

Advertiser	Approach	Creative pattern
Ingestible category leader	Ingestible, education-led	Explains the symptom as a gut problem, then sells a daily chew
Creator-led ingestible	Ingestible, creator-led	Partnership ads run from creator accounts, human-grade and vet-formulated claims, subscription pushed in the copy
Authority-persona vitamin	Authority persona	The same multivitamin creative delivered from a vet-named page
Topical incumbent	Category incumbent in the same active ingredient	Two-step relief kit, before-and-after results, multiple versions per ad
Direct challenger	Direct product comparison	Names three reasons other sprays failed: it stung, it rinsed off, the animal licked it off
Persona advertorial	Advertorial	Long first-person stories from individually named pages, attacking the prescription by name
Native-style testimonial	Native testimonial	Amateur photography, long copy, prescription-replacement claim

Structural observations

- **The problem is crowded and the solution is not.** Around 1,200 active ads speak to this symptom. Thirty-nine mention the active ingredient, and some of those are the client's own. The buyer searching for relief is sold an ingestible thirty times for every time they are shown a topical.
- **The loudest voices attack the prescription.** The advertorial pages lead with prescription drug names and promise an escape from them. That is effective and it is fragile. It invites policy action, and it cannot be said by a brand that wants veterinarians on its side.
- **Ingestibles sell a wait.** A chew or a multivitamin asks the owner to subscribe and watch for change over weeks. A topical works where the problem is, the same day. Nobody in the sweep is making that contrast.
- **The direct challenger is the one to watch.** It is the only advertiser in the sweep selling against other sprays on product mechanics, and its three objections are ones the client can answer directly.
- **Nobody pairs two credentials.** Vet-formulated and vet-approved are everywhere. The client's storefront states that every product is reviewed by both a veterinarian and a specialist. That pairing does not appear in any competitor ad we saw.

11. Unclaimed positioning

Five positions are available in this auction now, each supported by a finding above.

1. The second credential

Every competitor claims a vet. None claims a specialist. For a product whose entire promise is skin, the specialist is the more persuasive of the two, and it is already true.

2. Relief where the problem is

The category leader by ad volume sells something the animal eats. The client sells something that goes on the affected area. A direct, fair comparison of topical against ingestible, in the format the comparison static already proved, is open territory.

3. Safe on contact

The direct challenger has made "the animal licked it off" a named objection. The client's products are described as lick-safe, non-toxic and human-grade. An ad that shows the behaviour and explains why it is fine answers the objection in three seconds.

4. A routine, not a rescue product

The spray, rinse and wipes as a daily routine, in the way a person has a skincare routine. This is the position that raises order value and gives a reason to subscribe, and it connects directly to section 8.

5. Segment by segment

Each breed has its own version of the problem. The account tested these angles at small budgets. Buyers search and scroll that way, and almost nobody in the category addresses them accordingly.

12. 90-day remediation and scale plan

Days 1 to 7: stop the bleeding, fix the instruments

Nothing else in this plan is trustworthy until the measurement layer is fixed and spend is back at a level the account has shown it can carry.

- Reduce spend while the discrepancies are investigated.
- Audit the pixel and Conversions API event by event: add to cart, initiate checkout, view content and purchase. Add a third-party attribution tool so the numbers can be compared.
- Put every ad set on one agreed attribution setting, 7-day click for example, so every number from this point is comparable.
- Compare nine months of Meta purchases against store orders. Establish the new-customer share, the true ROAS, and confirm the targets.
- Relaunch assets B and C in a clean campaign with existing customers excluded.

Days 8 to 30: consolidate, restart the warm audiences

- Add new campaigns: one prospecting campaign with customer exclusions and a small number of broad ad sets, one creative testing campaign with real budgets per concept, and one warm campaign.
- Build three warm segments with their own creative and a frequency cap of three per week.
- Report on new-customer purchases alongside total purchases from day one of the new structure.
- Point prospecting traffic at the two-product set and the bundles, and test the free-shipping threshold as people buy them.
- Brief and produce the first batch of new creative: six to eight concepts.

Days 31 to 60: build the creative engine

- Launch new creative as a structured test against the relaunched controls, each with enough budget to reach a verdict.
- Add static and carousel formats to the new concepts.
- Test a first-order bundle above the \$60 threshold.
- Take the first scaling step, provided the gate conditions are met.
- Build and approve the peak-season slate.

Days 61 to 90: scale the budget

- Scale one step at a time, each gated on efficiency holding.
- Launch further concepts and retire fatigued assets on a defined schedule.
- Load and test the peak-season slate before CPMs peak.
- Compare 90-day performance against store orders, not only against Meta.

13. Budget scaling thresholds

The best sustained result of the year was a 2.10 ROAS in Q1 on roughly \$7,000 a month. The goal is to beat that figure on clean data and then carry it to a spend level the account has never held it at.

The ladder starts with a step down. That is deliberate. The last \$60,000 of quarterly spend returned 0.70. There is no version of this plan in which continuing to buy that spend makes sense while the structure is rebuilt.

Step	Monthly spend	Gate condition
0. Reset	\$18,000	Tracking reconciled with store orders. One attribution window. New-customer cost per purchase established. This gate must clear before any increase.
1	\$25,000	Account ROAS at or above 1.8 for 14 days. Warm campaign live and profitable. Monthly prospecting frequency below 2.5.
2	\$35,000	ROAS at or above 2.0 for 14 consecutive days. At least 8 distinct creative concepts live. Average order value at or above \$55.
3	\$50,000	ROAS at or above 2.1 sustained for 21 days. New-customer cost per purchase stable or falling. No single ad above 25% of spend.
4	\$70,000	ROAS at or above 2.1 at step 3 volume. Creative production sustaining 12 or more new assets a month. CPM increase under 20% against the step 0 baseline.
5	\$100,000	All step 4 conditions held for 30 days. Contribution margin confirmed positive at scale, including 90-day repeat revenue.

Conditions, not a timetable. All targets are provisional until the week one reconciliation and the margin figure are in hand.

Rules that govern the ladder

- No step increases spend by more than 50% at once. The April to September record in this account is what larger jumps look like.
- A minimum of 14 days at each step before the next is considered.
- Any step that fails its gate steps back one level rather than holding position.
- Frequency governs everything. If monthly prospecting frequency passes 3, the next dollar is going to people who have already seen the ad.

14. Creative concept territories

Each territory addresses a finding in this audit. Each is a territory rather than a single ad, so it supports multiple executions and can be tested and scaled on its own.

Territory 1: the proof shot

Open on the problem. Before, then after, in the first two seconds, filmed by a real owner. Product second.

Why it works: the in-house before-and-after is the most efficient asset in the account at 2.11 with a 66% to 73% hold rate. This territory makes more of what already works, with claim-safe copy.

Territory 2: the two credentials

The veterinarian and the specialist on camera, each explaining one thing in plain language. Asset B is the template.

Why it works: no competitor pairs the two credentials, and the vet-led asset returns 1.96 at the lowest cost per purchase in the table.

Territory 3: where the problem is

Topical against ingestible. A chew takes weeks. A spray goes on the spot. Statics and short comparison videos.

Why it works: the category's ad volume is ingestibles, and the one comparison static in the account is the only ad that explains the difference in a single frame.

Territory 4: the 30-second routine

Spray, rinse, wipes. After the walk, after the swim, before bed. Filmed as a routine.

Why it works: section 7 showed the order value problem. A routine sells three products where a rescue sells one.

Territory 5: safe on contact

The treated area is licked. The owner does not flinch. Then the ingredient list, in full.

Why it works: a competitor has made this a named objection. This answers it with something the client can show and the competitor has to explain.

Territory 6: segment by segment

One segment, one problem, one product per ad.

Why it works: these angles were tested at small budgets inside a campaign that could not fund them. They deserve a real test.

Territory 7: the refill

Built for past customers and warm audiences. Running low, the second product, the subscription.

Why it works: warm audiences are the cheapest revenue available and currently receive 0.13% of spend.

15. Creative production model

The plan requires 12 or more new assets a month by step four. That volume is where most ecommerce accounts stall, so it is worth being explicit about how it gets produced.

What we produce

- Concept development and scripting for all territories, including hook variants.
- Full creative briefs: shot lists, on-screen copy, timing, safe zones for each placement, and audio direction.
- Static and motion asset production, including re-edits of existing creator and in-house footage with new openings.
- All ad copy: primary text, headlines, descriptions and CTA variants, written to claim-safe language.
- Format adaptation across 9:16, 4:5 and 1:1.

What we need from you

- The raw files for assets B and C and every legacy asset, not the exported ads.
- The library of customer before-and-after photographs and reviews, with permissions.
- Usage rights and whitelisting status for every creator already under contract.
- New recording sessions if any assets are missing.
- Approval on brand voice and a written list of claims that have been cleared.

16. Seasonal planning

A 90-day plan starting in October runs directly into the most expensive auction of the year and ends just before the season when this product is needed most.

Peak season and Q4 gifting

- The August promotional sale is the cautionary example. \$1,120 spent, 9,933 people reached, a 3.04 frequency and a 1.30 return. The promotion was shown to a small pool that had mostly seen it already. A sale needs a warm audience built in advance and a prospecting slate that already works.
- Creative slate loaded and tested by 1 November. Testing new concepts during peak CPMs wastes money.
- Expect CPMs to rise 30% to 60% against the October baseline. The gates in section 13 should be read against that, not paused reflexively.
- The sets and the wider range give the brand a gifting offer that single bottles do not. Lead with them.

Allergy season

- The underlying problem peaks in spring and again in late summer, with a second seasonal driver in swimming season.
- The account's own history shows the cost of arriving late: CPM moved from \$35 to \$46 between March and April this year and stayed there.
- The ramp begins in February, with new concepts built and tested over the winter.
- Winter has its own angle: dry skin from indoor heating. It is a smaller season, and it is uncontested.

17. Projected outcomes, timelines and risks

These projections assume the week one measurement work is completed and the consolidated structure is launched as described. They are estimates, and every target is re-checked against reconciled data.

Window	Expected state	True ROAS
Days 1 to 14	Spend reset. Tracking audited. One attribution window. Reported ROAS likely falls as view-through credit is removed.	True baseline established
Days 15 to 30	Consolidated structure live. Legacy winners relaunched. Warm campaign restarted.	Baseline to 1.6
Days 31 to 60	New territories in test. Order value work live. First scaling step.	1.6 to 2.0
Days 61 to 90	Eight or more live concepts, warm audiences mature, peak-season slate tested.	2.1+

What will look worse before it looks better

Three things, all expected, all worth stating in advance so they are not misread as failure.

Attributed revenue will fall in the first two weeks because spend falls. At a 1.05 return, the revenue being given up is revenue that cost as much as it earned.

Reported ROAS will drop again when every ad set moves to 7-day click and reporting separates new customers from existing ones. The number on the dashboard will fall while the business is unchanged.

Consolidating campaigns resets the learning phase. Expect seven to ten days of unstable delivery after each structural change. That is why the plan sequences changes rather than making them all at once, and why a fourth rebuild without a measurement fix is unlikely to produce a different result from the third.

Risks

Risk	Likelihood	Mitigation
True new-customer ROAS is materially below 1.30	High	Expected. Targets re-baselined in week one. The plan's direction holds even if the starting point moves.
Margin and repeat rate do not support the target	Unknown	Resolved by the margin and 90-day value figures. If first-order economics are thin, the target becomes new-customer cost per purchase against customer value.
Tracking fixes are not implemented	Medium	Most fixes are configuration, not development. Each is scoped in week one, and any needing theme work is flagged.
Peak-season CPM inflation compresses efficiency	High	Seasonally adjusted gates. Slate tested before 1 November. Warm audiences built in October.
Creative claims trigger ad rejections	Medium	Claim-safe copy from the first brief, against a cleared-claims list agreed with the client.
Legacy winners have fatigued	Low to medium	They last ran at scale in June. Relaunch with new openings alongside the originals.

18. Reporting and working cadence

Weekly

- Monday scorecard: spend, ROAS, new-customer purchases, new-customer cost per purchase, average order value, prospecting frequency.
- What changed in the account last week.
- What is changing this week, and why.

Ongoing

- Optimisations and their effect on performance, with creative pipeline status: in brief, in production, in test, scaled, retired.

19. Summary of findings

- Spend rose 4.0x between Q1 and Q3. Revenue rose 2.0x. ROAS went from 2.10 to 1.05, and the added spend returned \$0.70 per dollar on Meta's own count.
- CPM rose 57% and the cost of reaching 1,000 people rose from \$55 to \$142. The budget was spent on frequency, not on new customers.
- Four different campaign structures ran in nine months. The one that returned 1.82 was switched off in July. The one that replaced it took 58% of the year's spend and returned 1.09 on a more generous attribution window.
- Purchases exceed initiated checkouts in every month since March, add-to-cart volume doubled in late August without checkouts moving, and content views dropped in late September. The mid-funnel cannot be read until that is fixed.
- The pixel already separates new customers from existing ones. No reporting or exclusion used it until September.
- The average order is \$49 and the Q3 cost per purchase is \$46. The business case depends on repeat purchase, and the order sits \$11 under the free-shipping line.
- The two most efficient ads of the year, both made in-house, are switched off. One creator carries 19.7% of spend, and the follow-ups to the lead ad returned half as much.
- The plan is to reduce budget deliberately, verify the numbers, and rebuild campaigns on a single attribution window with a different structure.

Prepared by Serafinity. Primary data pulled from the client's ad account and pixel dataset. Published with the client's written permission; identifiers redacted at their request.