Orpheus

AI operations for ecommerce agencies

For ecommerce creative agencies with 10+ people

Every workflow, in one list. *Pick the ones you want.*

14 creative jobs an AI system can take the first pass on. What each one replaces, what the system can do, and what your team still decides.

Workflow catalog · October 2026

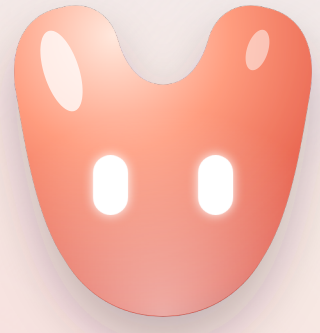

HOW TO USE THIS CATALOG

You don't need to know what to build. *Tick what hurts.*

Most agency owners only know what they want once they see the options. Read the list, tick the jobs that eat your team's week, and bring the numbers to the diagnostic.

STEP 1

Skim the list

Each workflow has a number. The next page shows all 14 by part of the week, then each one gets a card.

STEP 2

Tick three to five

Pick the jobs that cost your team the most hours or cause the most client complaints. You don't need to rank them.

STEP 3

Bring the numbers

Send them before your AI Ops Diagnostic. We measure each one in your agency and tell you which to build first.

These workflows describe what the system can do. Our own results come from email and retention work. Which ones fit your agency, and what they save, is measured against your own baseline in the diagnostic.

THE FULL LIST

14 workflows *across the creative week.*

RESEARCH

01 Customer research to ad angles
02 Competitor and category scan

BRIEFS AND CREATORS

03 First-draft creative briefs
04 Creator sourcing and footage tracking

PRODUCTION

05 Static-ad batches
06 Iteration shortlist
07 First cuts and hook variations

QA

08 Footage intake check
09 Edit QA and saved client rules

REPORTING

10 Weekly creative report per client
11 Creative variety check
12 Pre-launch test check
13 Ad-to-page check

CLIENT OPS

14 Morning list and client brain

01 RESEARCH

Customer research to ad angles

Today. A strategist reads hundreds of reviews and tickets by hand, or skips it.

With the system. Reviews, survey answers and support themes grouped into motivations, objections and confusion, each with a source quote. Then 3 to 5 original angles, each linked to the evidence behind it.

Your team decides. Which angles to make and which customer words can be used publicly.

Good first pick for any client with a decent review count.

02 RESEARCH

Competitor and category scan

Today. Someone scrolls ad libraries before every brief and keeps notes in their head.

With the system. A dated list of what competitors are running, which messages repeat, and gaps worth testing.

Your team decides. What's inspiration and what's off-limits. A competitor's ad running for months doesn't prove it's profitable.

Best with a short list of named competitors per client.

03 BRIEFS AND CREATORS

First-draft creative briefs

Today. Your best strategist writes every brief, or juniors write weaker ones.

With the system. A brief in your agency's template, pulled from the client's brand rules, approved claims, offer and past winners. It checks the footage, creators and editor time you actually have, and flags what's missing.

Your team decides. The final brief. A strategist edits instead of starting from blank.

Biggest lever when brief quality depends on who wrote it.

04 BRIEFS AND CREATORS

Creator sourcing and footage tracking

Today. A coordinator hunts for creators, chases footage and tracks usage rights in a spreadsheet.

With the system. A shortlist of creators with links and fit notes, draft outreach, and a tracker that flags overdue footage and missing rights.

Your team decides. Who to hire, every message sent, and all rates and terms.

Best with a busy UGC programme across several clients.

05 PRODUCTION

Static-ad batches

Today. Design time limits how many concepts each client gets to test.

With the system. A review board of 10 to 20 static ads across a few concepts, built from approved product images and copy. Product accuracy and text checked first.

Your team decides. Which ads go to the client. Designers polish the keepers.

Measure it by cost and time per approved ad, not ads generated.

06 PRODUCTION

Iteration shortlist

Today. Teams make brand-new concepts when a small edit to a working ad might do more.

With the system. A few existing ads worth another look, where each one loses people, and the smallest edit worth testing.

Your team decides. Which edits the footage supports. The media buyer owns the test.

Best with clients who already have spend and footage.

☐ 07 PRODUCTION

First cuts and hook variations

Today. Editors spend hours on rough cuts of talking-head and product footage.

With the system. A first cut from approved footage and script, bad takes removed, captions added, plus a few alternate hooks.

Your team decides. Pacing, continuity and the final edit. Complex motion work stays with editors.

Later pick. Trial it on simple formats first.

☐ 08 QA

Footage intake check

Today. Someone watches every clip to find out it's missing the shot you asked for.

With the system. Each delivery matched to its brief. Missing shots, missed lines, wrong product versions and unusable files flagged with timestamps.

Your team decides. What goes back to the creator. Unsure checks are marked for a person.

Good first pick if creators often send incomplete footage.

☐ 09 QA

Edit QA and saved client rules

Today. The client catches the typo, the wrong offer or the banned word. Then it happens again next month.

With the system. Edits checked for captions, spelling, offer wording, sizes and required elements. Every client correction is saved as a rule and checked on the next job.

Your team decides. The final sign-off before anything reaches the client.

Fewer revision rounds is the thing to measure.

☐ 10 REPORTING

Weekly creative report per client

Today. Someone rebuilds the same report every Monday, then still has to work out what to make next.

With the system. Results by angle and format, what changed since last week, what's still unclear, and a proposed next round of tests. Numbers come from the data, not the model.

Your team decides. The plan. The report owner checks numbers before it goes out.

Needs consistent ad names or IDs to work well.

☐ 11 REPORTING

Creative variety check

Today. A client gets 30 new ads a month that all say the same thing in different colours.

With the system. Recent ads grouped by message, customer situation and format, showing where spend is piling up and which angles haven't been tried.

Your team decides. Whether the groupings are right and what goes in the next batch.

Useful in quarterly reviews and renewal conversations.

☐ 12 REPORTING

Pre-launch test check

Today. A test goes live with two offers, three landing pages and a new discount mixed in. Afterwards nobody can say what worked.

With the system. Planned ads, offers and landing pages compared before launch. Mixed products, discount levels or page versions flagged, with a note on what the test can answer.

Your team decides. Which differences are intended. The media buyer owns the test design and budget.

Small pilot: one upcoming batch and a spreadsheet.

☐ 13 REPORTING

Ad-to-page check

Today. The ad promises one thing and the landing page talks about another. Nobody owns the gap.

With the system. Each ad compared with its exact landing page: product, offer, main promise and next step. A short list of mismatches and draft fixes.

Your team decides. What to suggest to the client. Page changes stay with whoever owns the site.

Easy add-on that shows clients you look past the ad.

☐ 14 CLIENT OPS

Morning list and client brain

Today. Client comments are spread across Slack, Frame.io, email and ClickUp. Brand rules live in one person's head.

With the system. Each person's day sorted before they start, from every channel. One place per client for products, approved claims, tone and past feedback, kept separate by client.

Your team decides. What gets done first. Account leads approve changes to client facts.

Foundation that makes every other workflow better.

YOUR SHORTLIST

The numbers you ticked

Write down the three to five that would help most, then send them when you book your AI Ops Diagnostic.

Book at orpheus.agency/book.html

WHERE TO START

Pick the job that hurts most. *Start with one.*

IF THIS SOUNDS LIKE YOUR TEAM	START WITH	A SMALL FIRST PILOT
Research and briefs take days	Customer research to ad angles	One client, one product: a sourced research brief and 3 to 5 angles.
We can't make enough variations to test	Static-ad batches	One product, a review board of 10 to 20 ads. Measure time per approved ad.
Footage and edits come back wrong	Footage intake check	The next creator deliveries for one client, checked against their briefs.
Clients repeat the same feedback	Edit QA and saved rules	One client's past feedback turned into rules, then every new edit checked.
Reports eat a day a week	Weekly creative report	One client, a past period checked against your report, then three live weeks.
Quality depends on who's on the account	First-draft creative briefs	Briefs for one client, side by side with your senior strategist's own.
Our tests don't tell us much	Pre-launch test check	The next planned batch for one client, checked in a spreadsheet.

HOW WE'D START

One workflow first. *Measured against today.*

STEP 1

AI Ops Diagnostic

We map where your creative team's hours go and which jobs repeat across clients. You keep the scorecard and plan whether or not you hire us.

STEP 2

Pick and baseline

Choose the workflow with the clearest payoff. Record how long it takes today and what good looks like.

STEP 3

Build and test

Set it up in your tools and run it on real client work, with your team reviewing every output.

STEP 4

Expand and reprice

Add the next workflows, train the team, and rebuild one service line so the freed time is sold as bigger work.

Book an AI Ops Diagnostic →

orpheus.agency/book.html