


A FIVE-DAY BUILD GUIDE



# THE COHORT LTV DASHBOARD.

Six numbers most brands have never put on a single page — and the operating manual for putting them there. Built once, refreshed monthly, used to decide where every next dollar goes.

---

FOR  
DTC Operators  
& Finance Leads

EFFORT  
5 working days

OUTPUTS  
1 dashboard,  
3 decisions/month

PAGES  
01 / 08

# CONTENTS & THE THESIS IN ONE PARAGRAPH.

Cohort LTV is the per-monthly-cohort cumulative profit your customers generate over time. **Blended LTV cannot answer the three questions that decide where your next ad dollar goes.** This guide walks the build, end to end.

§ 01	<b>The skill, in three questions</b>	P. 03
	What blended LTV hides, and why cohorts don't.	
§ 02	<b>Five mistakes that show up in every first audit</b>	P. 04
	The patterns that produce a confident, wrong number.	
§ 03	<b>The six numbers, per cohort</b>	P. 05
	Definitions, formulas, and what each one decides.	
§ 04	<b>The five-day build plan</b>	P. 06
	Day-by-day tasks, owners, and deliverables.	
§ 05	<b>Reading the dashboard &amp; making the call</b>	P. 07
	The three decisions you make every month, and how.	
§ 06	<b>The reallocation, illustrated</b>	P. 08
	What we typically find in the first 60 days.	

## Who this is for

Founders, finance leads, and growth operators at DTC brands doing \$5M–\$200M in annual revenue who have a working data warehouse, a paid-media program with two or more sources, and the suspicion that blended LTV is lying to them.

## What you'll need

Order-level export (customer ID, date, gross, COGS, fulfillment, discounts, refunds), paid-spend by source by day, a first-touch or last-click attribution table, and a query tool — Looker, Mode, Hex, or a spreadsheet for the v1.

# THREE QUESTIONS A BLENDED LTV NUMBER CANNOT ANSWER.

Cohort LTV is cumulative profit by month of first purchase, plotted over time. Once you have the curve — by cohort, by acquisition source — three decisions become legible that were invisible before.

## 01 · Are this month's customers more or less valuable than last month's?

A blended LTV averages every customer ever acquired. It cannot tell you whether the cohort you bought in February is on a steeper or flatter curve than November's. Without that signal you cannot tell if your product, your offer, or your channel mix is improving.

## 02 · Which channel produces the highest 12-month value, not just the cheapest first order?

Paid social is often the cheapest first purchase and the lowest twelve-month contribution. Email and organic search are typically the opposite. A blended pool averages these into a number that matches neither and recommends nothing.

## 03 · How fast does a cohort pay back its CAC?

Payback period — the day index where cumulative contribution dollars cross CAC — is what determines how aggressively you can scale. LTV without payback timing is a vanity number; payback without LTV is a constraint without a ceiling. You need both.

Without this dashboard, brands optimise CAC against a number that does not exist — and discover too late that the cheapest customer to acquire was the worst one to keep.




FIG. 01 — CUMULATIVE CONTRIBUTION DOLLARS VS. CAC, BY SOURCE HYPOTHETICAL COHORT, ILLUSTRATIVE

# FIVE MISTAKES THAT SHOW UP IN EVERY FIRST AUDIT.

Each one produces a number that looks correct, gets quoted in board meetings, and quietly funds the wrong channel for another quarter.

---

**M / 01**     **Computing one blended LTV across all customers, then setting CAC targets against it.**

The blend hides the difference between a \$400 LTV email subscriber and a \$90 LTV paid-social one-and-done buyer. Targets set against the average overpay for the bad customer and underpay for the good one.

---

**M / 02**     **Counting gross revenue per cohort instead of contribution dollars.**

Returns, refunds, and merchandise discounts are 10–25% of revenue for most DTC brands. The cohort that looks healthy on revenue can be 40% thinner on contribution — and contribution is what funds the next quarter.

---

**M / 03**     **Measuring LTV at intervals that are too short to be honest.**

A 90-day LTV ignores second-purchase behaviour that often does not show up until day 120. Use D0 / D30 / D60 / D90 / D180 / D365 — and read the slope between D90 and D180 as the leading indicator.

---

**M / 04**     **Not segmenting by acquisition source.**

Organic search, paid social, email, and affiliate produce wildly different LTV curves. Treating them as a blended pool is the most expensive analytical shortcut in DTC.

---

**M / 05**     **Looking at LTV without a CAC payback period attached.**

LTV without payback timing is a brag number. LTV minus CAC over time is the operating number. Always ship them together.

---

The throughline: every shortcut produces a confident, wrong number. Cohorts, contribution, full intervals, source segmentation, and payback are not optional refinements — they are the dashboard.

# THE SIX NUMBERS, PER COHORT.

Build these once. Refresh monthly. Read them together.

## 01 Cohort definition

**DEFINITION** Group customers by month of first purchase. The minimum useful cohort size is 200 customers; below that, sample noise eats the signal. If a monthly cohort is smaller, roll into quarterly cohorts until the math stabilises.

```
cohort = MONTH(first_order_date), n ≥ 200
```

## 02 Cumulative gross revenue per cohort

**TOP LINE** Sum every order tied to every customer in the cohort across days 0–30, 60, 90, 180, 365. Plot the curve. The slope tells you whether your customers come back, and how fast.

```
Σ orders.gross by cohort × {30, 60, 90, 180, 365}
```

## 03 Cumulative contribution dollars per cohort

**CASH LINE** Same curve, after subtracting COGS, fulfillment cost, returns cost, and merchandise discounts. This is the line that matters for cash — and it's the one most brands have never plotted.

```
contrib = gross - COGS - fulfillment - returns - discounts
```

## 04 CAC by acquisition source per cohort

**ACQUISITION** Pull paid spend, attribute by source, divide by new customers acquired in that source for that cohort. Last-click attribution is fine for the v1; multi-touch is the v2.

```
CAC_src = spend_src / new_customers_src
```

## 05 Payback period — by cohort, by source

**TIMING** The day index where cumulative contribution dollars equal CAC. Sub-90-day payback: scale aggressively. 180-day: scale carefully. Over 365-day: the channel needs a different model, or it needs to be cut.

```
payback_d = min{ d : cum_contrib(d) ≥ CAC }
```

## 06 LTV : CAC ratio by source at 12 months

**THE RATIO** The 3:1 rule applies. Below 2:1 the source is losing money on a cash basis once you load in storage and overhead. Above 4:1 the source is underfunded — move budget toward it.

```
LTV:CAC = cum_contrib(365) / CAC
```

# FIVE WORKING DAYS, END TO END.

One analyst, one finance partner, one growth lead. Source data assumed accessible. Each day ships a deliverable the next day depends on.

DAY

01

## Define cohorts & pull the order spine

Lock the cohort grain (monthly, fallback quarterly under  $n=200$ ), confirm timezone, and pull the order-level table that everything else joins to.

- Customer ID, first-order date, cohort label, acquisition source
- Order-level: gross, COGS, fulfillment, discount, refund, return flag
- Cohort-size sanity check & quarterly rollup rules

→ → **ORDERS\_BY\_COHORT TABLE, VALIDATED**

DAY

02

## Build cumulative gross & contribution curves

From the order spine, compute cumulative gross and contribution per cohort at D30 / D60 / D90 / D180 / D365. Plot both — the gap between them is your hidden margin story.

- Cumulative SQL with a day-index window function
- Reconcile against the P&L within  $\pm 2\%$
- Save as cohort\_curves

→ → **NUMBERS 02 & 03 SHIPPED**

DAY

03

## Attribute CAC by source, by cohort

Pull paid-spend by source by day. Map customers to source via last-click for v1. Compute CAC per source per cohort. Flag any source with fewer than 50 customers in a cohort as low-confidence.

- Spend table joined to customer table on source + week
- Last-click for v1 — note the assumption in the dashboard footer
- Validate spend totals against the ad-platform invoice

→ → **NUMBER 04 SHIPPED**

DAY

04

## Compute payback & LTV:CAC

Join the contribution curve to CAC. Solve for payback day index per cohort  $\times$  source. Compute the 12-month LTV:CAC ratio. Bucket each source as *scale / hold / fix or cut*.

- Payback bands: <90 / 90–180 / 180–365 / >365
- LTV:CAC bands: <2 / 2–3 / 3–4 / >4
- Confidence flags on cohorts still maturing

→ → **NUMBERS 05 & 06 SHIPPED**

DAY

05

## Wire the dashboard & the monthly refresh

One page. Six numbers. Filters for cohort and source. A monthly scheduled refresh and a 30-minute review cadence on the books.

- Cohort heatmap, contribution curve, payback bar, LTV:CAC table
- Footer with attribution model, data freshness, last refresh time
- Calendar invite: monthly review, 30 minutes, three decisions

→ → **DASHBOARD LIVE, REFRESH SCHEDULED**

# BUILT ONCE. REFRESHED MONTHLY. USED TO MAKE THREE DECISIONS.

The dashboard is not a report. It is a forcing function for monthly capital allocation.

## 01

### Where to add ad spend

Sources with LTV:CAC above 4:1 and payback under 180 days. They are underfunded. Move budget toward them until the ratio compresses to ~3:1 or payback drifts past 180 days.

## 02

### Where to cut ad spend

Sources with LTV:CAC under 2:1 and payback over 365 days. They lose money on a cash basis once storage and overhead are loaded in. Cut, redesign the offer, or both.

## 03

### Which SKUs & offers improve repeat rates

Inside the cohorts that matter most, look at second-order SKU mix. The repeat-rate winners get merchandised harder; the one-and-done SKUs get re-priced or sunset.

### The monthly review, in thirty minutes

**00–05 min** New cohort lands. Compare its D30 contribution slope to the trailing three-month average. Flag drift beyond  $\pm 10\%$ .

**05–15 min** Walk the LTV:CAC table by source. Mark every row as *scale / hold / fix or cut*. No row is allowed to stay unmarked.

**15–25 min** Reallocate. The committee can move up to 20% of monthly spend without escalation. More than that, escalate.

**25–30 min** Write the three decisions in the dashboard footer with date and owner. Adjourn.

A dashboard that does not produce a decision is overhead. Three decisions a month, written down, with an owner — that is the unit of value.

# WHAT THE FIRST 60 DAYS USUALLY REVEALS.

We build this dashboard for every Anata client in their first 60 days. The most common discovery has the same shape across categories: paid social CAC is materially lower than the brand thought — **and so is its LTV.**

SOURCE	CAC, PERCEIVED	CAC, ACTUAL	12-MO LTV	LTV : CAC	VERDICT
Paid social	\$48	\$36	\$62	1.7×	Sub-2:1 — <i>fix or cut</i>
Email / lifecycle	\$22	\$14	\$71	5.1×	Underfunded — <i>scale</i>
Organic search	\$18	\$11	\$58	5.3×	Underfunded — <i>scale</i>
Affiliate	\$31	\$29	\$74	2.6×	In-band — <i>hold</i>

Email and organic are sub-1:1 on CAC and 5:1 on LTV — but they are starving for budget because the brand has been chasing top-line first orders. The reallocation typically moves **15–25% of monthly ad spend** into retention, lifecycle, and SEO.

## What changes in six months

Within six months, gross margin lifts 2–4 points and revenue compounds because the same dollar buys a customer worth roughly twice as much over twelve months. Contribution per new customer is the metric that's moved; the rest follows.

The dashboard does not require a re-platform, a new attribution stack, or a consultant on retainer. It requires the discipline to read six numbers per cohort, every month, and to write down three decisions that follow from them.

Build it once. Refresh it monthly. Let it embarrass you the first time. Then let it earn its keep every month after.

**End.** Anata · Field Guide № 04 — *The Cohort LTV Dashboard*. Compiled April 2026. For internal & client use. The hypothetical figures on this page are illustrative; replace with your own data on first build.