OPERATOR-LED GROWTH

The Customer Acquisition Cost Field Guide

Why the benchmark you are using describes a company that is not yours, and how to compute a number you can defend.

Growth Revenue Optimization

Every figure sourced. 2026 edition.

ABOUT GRO

Growth Revenue Optimization

gRO runs a company's entire growth function through one senior operator: paid acquisition, lifecycle and email, positioning, performance creative, conversion optimisation, and analytics. One person owns both the strategy and the hands-on execution, works inside your own ad accounts and CRM, and answers to the same pipeline number you read. An AI agent fleet handles the production work underneath, so one senior person can run the whole function.

Ad spend flows through your accounts, not ours. No markup, no media commission. The retainer is the only fee. For B2B and B2C companies between $1M and $10M in revenue.

The firm was founded by Ro Maldonado, who has spent fifteen years inside funnels across B2B SaaS, fintech, wealth management, retail banking, and consumer brands.

ABOUT THIS GUIDE

Every statistic in this document was verified against a named, dated primary source that was actually opened and read. Sample sizes are printed at the metric level rather than the headline level, because the two are rarely the same. Where credible sources disagree, the disagreement is printed rather than resolved in favour of the more convenient number.

A short list of widely circulated figures was considered and cut for lack of a sourceable original. That list is printed too, on page 19, because what a document refuses to tell you is part of what it is worth.

This guide contains no claim about what gRO would achieve for your company. It is a method and a set of benchmarks. Any number describing gRO's own work, or Ro's record before founding it, is labelled as such where it appears.

Contents

IN BRIEF

The number you are benchmarking against describes someone else.

Ask what a company should spend to acquire a customer and you get four credible answers, all published in the same eighteen months, all from named research bodies with stated methods. Marketing runs at **7.8% of revenue**[1] or **9.0%**[2] or **13.3%**[2] depending on who was asked, and a B2B SaaS company between $1M and $5M in ARR spends **35% of revenue** on sales and marketing combined.[3] These are not competing estimates of one quantity. They are four different quantities wearing one label.

The gap matters most at the bottom. The famous number, the one that circulates in board decks, comes from a sample that is mostly billion-dollar companies. Asked of firms under $10M in revenue, the same survey returns a figure **nearly double** the one reported by companies twenty times their size.[2]

Meanwhile the confidence is not in doubt. **84% of marketers** say they are extremely or very confident in their ability to measure return on marketing investment. **38%** measure across their traditional and digital spend together.[4] Those two numbers describe the same people.

This guide is about closing that distance. It sets out what a defensible acquisition-cost number contains, which customers belong in the denominator, the five places the calculation reliably breaks, and what the published benchmarks do and do not entitle you to conclude. It ends with a one-page worksheet. If you cannot fill in three of its boxes from a system of record, that is the finding, and it is a more useful one than the number at the bottom.

What the benchmarks actually say

Published acquisition benchmarks are real, rigorous, and mostly about companies that are not yours. Read correctly they are still the best instrument available.

By Ro Maldonado

At a Glance

- Four credible sources answer "what share of revenue goes to marketing" with **7.8%, 9.0%, 13.3% and 35%**. The spread is denominators and populations, not disagreement about the world.
- For a company between $1M and $10M, the two best-sampled sources differ by roughly **two times** on the same nominal measure. Benchmarking against the wrong one will tell a healthy company it is badly underinvesting.
- The ratio of lifetime value to acquisition cost is **worst in the $1M to $5M band**, below both smaller and larger companies. If it feels harder at your size than it looked from below, the data agrees with you.

The question sounds simple enough to have one answer. It does not. Below are four figures published inside eighteen months, each by a named body with a stated method, each answering some version of "what share of revenue goes to acquiring customers." They range across a factor of four.

Nothing here is wrong. The spread is produced by three things: what sits in the numerator, whether the denominator is revenue or ARR, and above all who was asked. The fourth bar counts sales *and* marketing, so it is not measuring the same quantity as the first three and is drawn differently to say so.

Figure 1: One question about marketing spend has four credible answers, and the largest is more than four times the smallest

Marketing spend as a share of revenue, by source and sampled population (%)

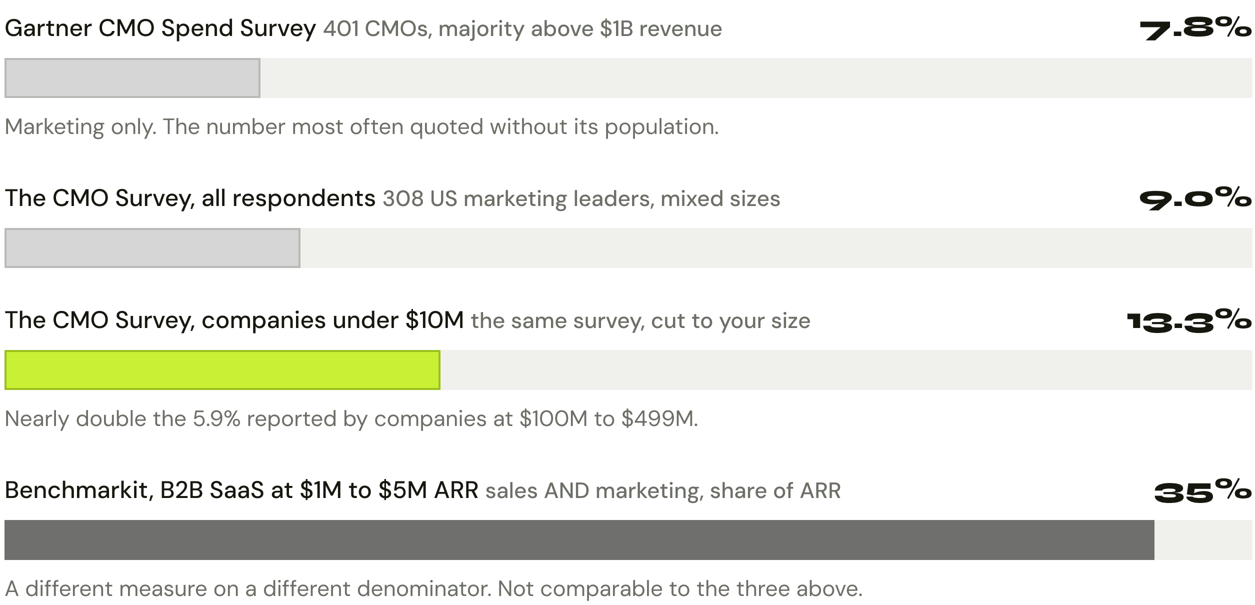

Note: The first three count marketing spend against revenue. The fourth counts sales and marketing combined against ARR, and its interquartile range is 20% to 60%, n=157. Gartner's sample is majority billion-dollar companies and does not describe a company below $10M
Sources: Gartner 2026 CMO Spend Survey[1]; The CMO Survey, February 2026[2]; Benchmarkit 2025[3]

Narrow the question to companies the size of yours and the disagreement does not resolve. It sharpens. Two of the best-sampled sources available, both surveying private B2B SaaS companies, both reporting sales and marketing as a share of the top line, differ by roughly a factor of two.

The explanation is not that one is careless. It is that they sampled different companies. Benchmarkit's respondents skew venture-backed; SaaS Capital's skew bootstrapped. Equity-backed companies spend **100% more on marketing** and 70% more on sales than bootstrapped peers, and **83% of bootstrapped companies** sit at or near breakeven against 52% of equity-backed ones.[5] They are running different businesses with different costs of capital, and averaging them describes neither.

Figure 2: For a company at your size the same nominal measure differs roughly two-fold, and the difference is who was sampled

Sales and marketing spend as a share of the top line, private B2B SaaS below $20M (%)

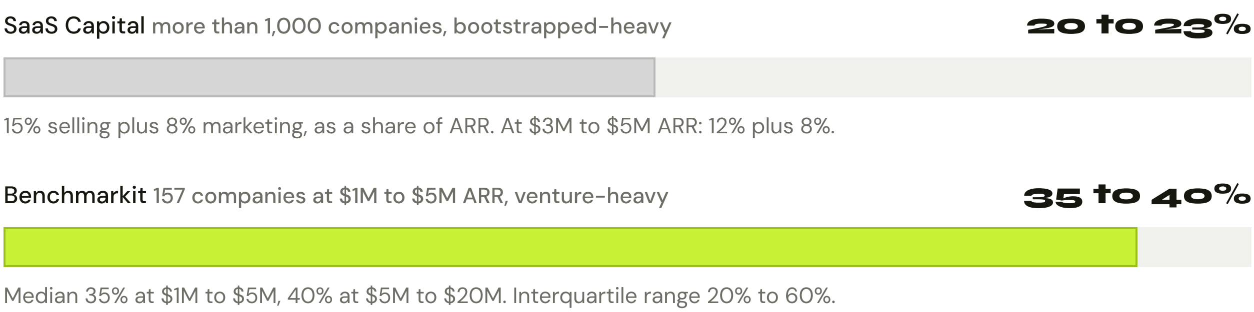

What this costs you if you get it wrong. A bootstrapped company spending 22% of revenue on sales and marketing, benchmarking itself against the 35% figure, will conclude it is underinvesting by a third and raise spend accordingly. On the other source it was already at the median.

Note: Denominators differ: SaaS Capital reports against ARR, and its survey completed March 2026. Benchmarkit's own data shows venture-backed companies at 45% to 47% against private-equity-backed at 33%, which is most of the gap
Sources: SaaS Capital 2026[5]; Benchmarkit 2025[3]

Two figures worth holding together. The first says the economics of acquisition are hardest in a specific revenue band, and it is probably yours. The second says the time it takes to earn a customer back has been getting longer, and that the spread around the middle is wide enough to make the median nearly useless on its own.

Figure 3: The return on acquisition spend is worst in the $1M to $5M band, below both smaller and larger companies

Median ratio of lifetime value to customer acquisition cost, by ARR band (x)

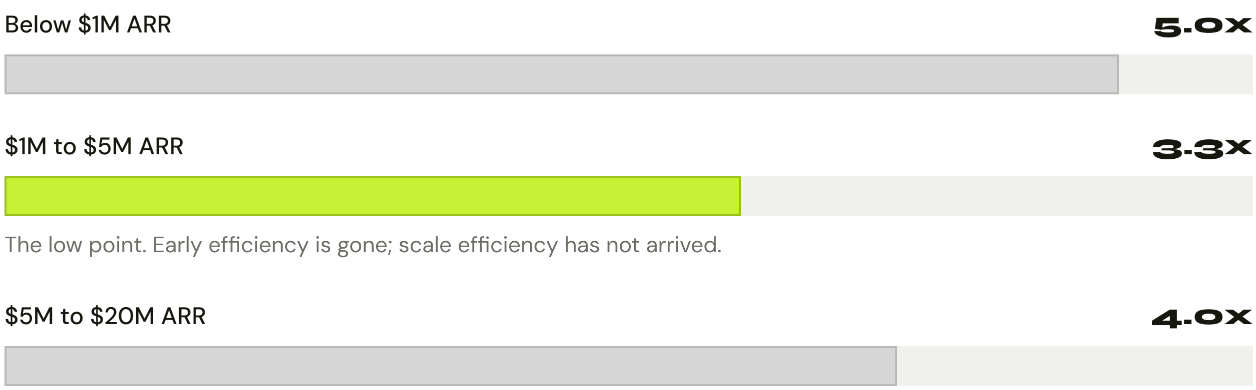

Note: n=101. A ratio below 3.0 is conventionally read as unsustainable, though the convention itself is a heuristic rather than a measured threshold
Source: Benchmarkit 2025[3]

Figure 4: Time to earn back an acquired customer has lengthened, and the spread around the median is twice the median itself

Months to recover customer acquisition cost, B2B SaaS (months)

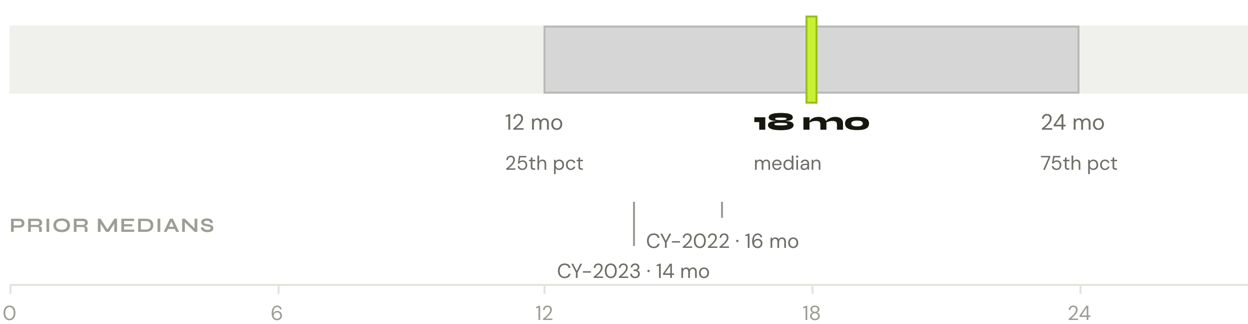

The spread is the finding, not the median. Half of all companies sit outside 12 to 24 months, so a twelve-month rule of thumb describes the top quartile rather than the middle. And the median has not moved in one direction: 16 months, then 14, then 18.

Note: n=148 for CY-2024. A later cut of CY-2025 data from the same research house reports the median improving to 16 months on a different cohort of 198 companies, so treat the direction as unsettled rather than reversed
Sources: Benchmarkit 2025[3]; Aleph and Benchmarkit 2026[6]

The arithmetic

Almost nobody disputes the formula. Both halves of it are ambiguous, which is how two honest people compute the same company's acquisition cost and land a factor of three apart.

By Ro Maldonado

At a Glance

- Six cost lines belong in the numerator and are routinely left out of it. The most commonly omitted is the one most likely to be large: **the agency retainer or management fee**.
- There are three defensible denominators and they are not interchangeable. Blended is always the most flattering, because it takes credit for customers who arrived on their own.
- A cost figure with no stated scope is not a number, it is a category. Write it as $\mathbf{CAC}_{\text{paid}}$ or $\mathbf{CAC}_{\text{blended}}$ or do not write it at all.

The formula is total acquisition cost divided by customers acquired. The first argument is about what "acquisition cost" contains. Every exclusion below is individually defensible, and a company that makes all of them at once will report a number roughly half its real one.

The test that settles it: take the fully-loaded figure to whoever owns the P&L rather than to whoever owns the campaigns, and ask what is missing.

Figure 5: Six cost lines belong in the numerator, and the ones most often excluded are the ones most likely to be large

Cost lines, whether they belong in customer acquisition cost, and why each is contested

COST LINE	INCLUDE	WHY IT IS ARGUED ABOUT
Paid media spend	Yes	The one nobody disputes.
Agency or contractor fees	Yes	A cost of acquiring the customer. Excluding it is the single most common way a flattering number is produced, and at this company size the fee is often a third of the media.
Marketing salaries, loaded	Yes	A team that costs more than the media it buys is not unusual below $10M.
Creative production	Yes	Video, design and copy are acquisition costs even when they are booked to brand.
Martech and tooling	Yes	Attribution, email platform, and the CRM seats used by acquisition.
Sales commission on new logos	Yes	Left out, the result stops being a cost-to-acquire figure.
Sales base salaries	Judgment	Include where the role exists to close new business. State which way you went.
Customer success and onboarding	No	Cost to serve, not cost to acquire.
Expansion and upsell effort	No	Belongs to net revenue retention.
Brand spend with no acquisition intent	Disclose	Defensible to exclude. Not defensible to exclude silently.

Note: This table is gRO's method, not a published standard. No research body has issued an authoritative inclusion list, which is part of why published acquisition-cost benchmarks are difficult to compare
Source: gRO analysis

The second argument is about who counts as an acquired customer. Three answers are defensible. They are not interchangeable, and the difference between them is not small.

All new customers gives you a blended figure that includes everyone who arrived organically, by referral, or by word of mouth. It is always the most flattering. **Paid-sourced customers only** tells you whether the paid channel works. **New logos only** excludes reactivations, on the ground that a returning churned customer is not an acquisition.

The illustration below runs one identical month of spend through all three. Nothing changes except the denominator.

Figure 6: One month of spend produces three defensible acquisition costs that differ by a factor of nearly three

Illustrative: $60,000 fully-loaded spend, 120 new customers, of which 55 paid-sourced and 8 reactivations ($)

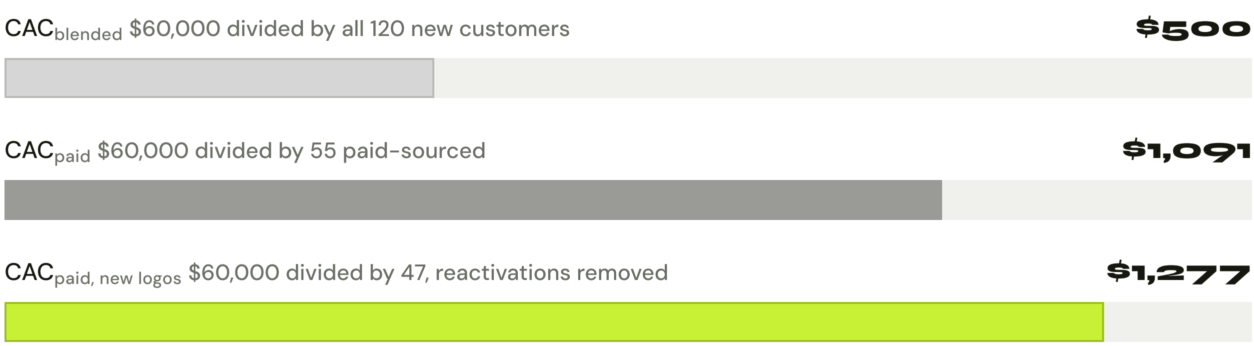

The most conservative reading, and the one that answers "what does the next customer cost".

Note: Arithmetic on stated inputs, not a measurement of any company. The inputs are chosen to be plausible and are not drawn from an engagement
Source: gRO illustration

The rule that makes it honest. State the numerator and the denominator on the same line as the number. An unqualified acquisition cost is a category, not a measurement.

Where the number goes wrong

Five failure modes, in the order they cost the most. Each comes with a test you can run yourself, because a leak you find is more convincing than a leak you are told about.

By Ro Maldonado

At a Glance

- The most expensive error is using a blended figure to make paid decisions. Organic volume subsidises the average, so every additional dollar buys worse than the average implies.
- Add up what every advertising platform claims it produced. If the total exceeds the customers you actually gained, the excess is **double counting**, and it is almost always there.
- A good acquisition cost with an unexamined payback period is a financing problem wearing a marketing costume.

01 A blended figure standing in for a paid one

Acquisition cost looks acceptable, but raising spend does not raise customers in proportion. Organic, referral and word-of-mouth volume is subsidising the average, so the marginal dollar performs materially worse than the average suggests. This is the most expensive error on the list because it survives scrutiny: the number is real, it is simply answering a question nobody asked.

> **The test.** Compute both. If $\textbf{CAC}_{\text{blended}}$ and $\textbf{CAC}_{\text{paid}}$ differ by more than about a third, every decision made on the blended figure was made on the wrong one.

02 An attribution window that flatters the channel

Long click-through and view-through windows let a channel claim conversions it merely witnessed. The error compounds because each platform is measuring in its own walled garden and none of them net off against each other. Meanwhile the underlying cost of the channel is moving: Meta's average price per ad rose **9% in 2025**,[7] and on Google, brand-keyword costs per click rose **19% year over year** against 3% for non-brand[8] which means a rising share of paid budget is being spent defending terms the company already owns.

> **The test.** Add up the conversions every platform claims for a single period. Compare the total to the new customers in the CRM. Ask nobody to explain the gap until you have measured it.

03 Costs that live outside the marketing line

The retainer booked to operations, the creative booked to brand, the tooling booked to IT, the commission booked to sales. Each exclusion is defensible on its own and collectively they halve the number. This is rarely deception. It is org-chart accounting meeting a metric that does not respect org charts.

> **The test.** Take the fully-loaded figure to the person who owns the P&L, not the person who owns the campaigns, and ask what is missing.

04 Cohort mixing

An acquisition cost computed across a period in which the mix changed, a channel launched, a discount ran, a segment opened, is an average of two different businesses. The average describes a company that does not exist, and it is the company you will then forecast.

> **The test.** Cut the same period by month and by channel. If the monthly figures do not sit close to the average, the average is describing nothing real.

05 A good number with a payback nobody checked

Acquisition cost is a cost, not a verdict. The same figure is excellent at a four-month payback and insolvent at twenty-four, because the second is a financing problem wearing a marketing costume. The published median has been drifting the wrong way, and the spread around it is wider than the median itself.[3]

> **The test.** Divide acquisition cost by monthly gross profit per customer. If the result exceeds your runway, acquisition cost is not the number that matters this quarter.

None of these five requires a tool you do not have. Four of them are arithmetic on data already sitting in an ad account and a CRM. The constraint is almost never the instrumentation. It is that nobody has been given the afternoon, and the standing to ask the finance team an awkward question.

There is one more reason a paid acquisition cost can be honest and still be worse than it looks. Some of the money never arrives. The most rigorous study of this ran log-level data on $123M of programmatic spend across 35.5 billion impressions, and found that **36 cents of the advertiser's dollar** reached the consumer.[9]

Two things about that figure deserve emphasis. Agency fees were explicitly outside its scope, so the real fraction is lower. And the loss is not fraud: the same programme's later work puts median invalid traffic at **0.5% of spend**, a rounding error next to the 35% lost to impressions that were never viewable.[10] The money is not being stolen. It is being spent on advertising nobody saw.

Figure 7: Roughly a third of a programmatic dollar reaches the consumer, and fraud is not where the rest of it goes

Disposition of one dollar of programmatic media spend entering a demand-side platform (cents)

Supply-chain transaction costs
the technology take between advertiser and publisher

Lost media productivity
non-viewable, non-measurable, made-for-advertising sites

Reaches the consumer
and agency fees were outside the study, so the true figure is lower

Fraud is not the leak. Median invalid traffic measures **0.5% of spend**. The losses are non-viewable impressions (35.2%) and non-measurable ones (13.3%). Headlines about tens of billions lost to ad fraud are contradicted by the only industry dataset with log-level impression data.

Notes: Figures describe programmatic display bought through demand-side platforms and do not transfer to search, or to social bought directly. The three components are drawn from the 2023 supply-chain study; the media-quality decomposition is from the later benchmark programme
Sources: ANA Programmatic Media Supply Chain Transparency Study, 2023[9]; ANA, TAG TrustNet and Fiducia benchmarks, 2025 and 2026[10]

Underneath all five failure modes sits a single uncomfortable pattern. The confidence that marketing measurement is under control substantially exceeds the evidence that it is, and both have been measured on the same populations.

This is not a case against measurement. It is the case for auditing the measurement before acting on it, and it is why this guide ends with a worksheet rather than a recommendation.

Figure 8: Confidence in marketing measurement runs far ahead of the capability behind it

Share of marketers reporting each, most recent available reading (%)

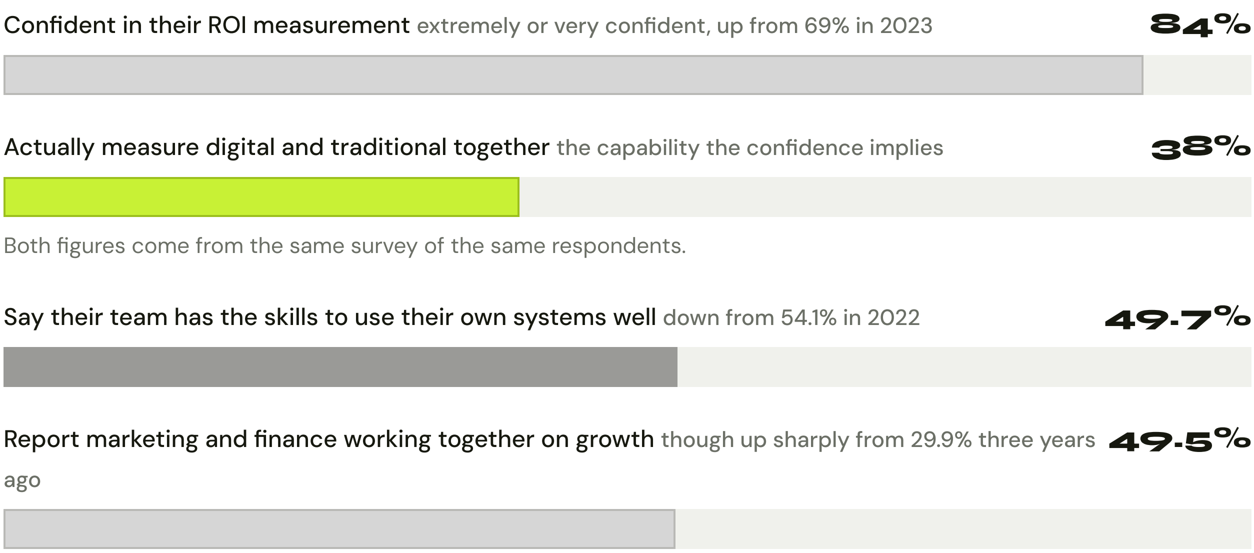

Note: The first two are from a survey of approximately 2,000 global marketers managing budgets of $1M or more, which skews larger than this guide's audience. The second two are from a US cross-section in which 97% of respondents are VP or above
Sources: Nielsen 2024 Annual Marketing Report[4]; The CMO Survey, February 2026[2]

One further finding belongs beside these. Acquisition budgets are now **26.0% larger** than retention budgets, up from 19.6%, in the same period that retention performance grew 12.8% against acquisition's 7.4%.[2] The survey reports both without reconciling them. When profits miss, marketing is the line that gets cut **45.4% of the time**.[2]

The worksheet

Fill this in from systems of record only. Where you have to ask someone, or estimate, mark the box rather than filling it. The count of marked boxes is the finding.

Figure 9: A defensible acquisition cost needs sixteen inputs, and the ones you cannot source are the result

Worksheet: fully-loaded customer acquisition cost and payback, one period

	INPUT	VALUE
A	Paid media spend, this period	
B	Agency, contractor and management fees	
C	Marketing salaries, fully loaded	
D	Creative production	
E	Martech and tooling	
F	Sales commission on new logos	
G	**Total acquisition cost** **A plus B plus C plus D plus E plus F**	
H	New customers, all sources	
I	New customers, paid-sourced	
J	Of I, reactivations to remove	
K	**New paid logos** **I minus J**	
L	$\text{CAC}_{\text{blended}}$ G divided by H	
M	**$\text{CAC}_{\text{paid, new logos}}$** **G divided by K**	
N	Monthly gross profit per customer	
O	**Payback, in months** **M divided by N**	
P	**Boxes you could not fill from a system of record** **the finding**	

Note: Gross profit, not revenue, at line N. Using revenue overstates payback by the whole of cost of goods sold, which for a services-heavy company is not a small error
Source: gRO method

What to do with the answer

Three outcomes. None of them is a reason to call anyone.

You filled every box and the payback sits inside your runway.

Nothing in this guide is your problem. The constraint is elsewhere, and the next honest question is which channel is scaling sub-linearly, which is a different investigation with different instruments.

You filled every box and the payback sits outside your runway.

This is a pricing, margin or retention question wearing a marketing costume. Cutting spend will improve the ratio and shrink the business, which is why it is so often the first move and so rarely the right one. The lever is gross profit per customer, at line N, not the numerator above it.

You could not fill three or more boxes.

The number is not wrong. It is absent, and every decision made against it has been made against a guess with a decimal point on it. Fix the measurement before changing the spend. This is the most common of the three outcomes by a wide margin, and it is the one nobody puts in a board deck.

If you would rather not do it yourself

Two ways, and the first one costs nothing. gRO publishes a free 47-point Operator's Funnel Scorecard at **applygro.com/audit**: seven sections, 47 checks, about twenty-five minutes, no email gate and no call. It covers this worksheet and the rest of the funnel around it.

If you want the work done rather than the checklist, the engagement that fits is a **Funnel Audit, $1,500, one week**: a written diagnosis of the funnel with prioritised fixes, read out of your own ad accounts, CRM and attribution. It is a standalone piece of work and it is designed to be useful whether or not anything follows it.

A guide is partly defined by what it declines to tell you. Each figure below is in wide circulation and would have strengthened this document's argument. Each was traced to its origin and cut. They are printed here so that the ones that survived can be read as having survived something.

Figure 10: Six widely circulated figures were considered for this guide and cut, each for a specific reason

Rejected claims, the argument each would have supported, and the reason for exclusion

CLAIM	WOULD HAVE SUPPORTED	WHY IT WAS CUT
"Spend 7 to 8% of revenue on marketing" attributed to the US Small Business Administration	A simple benchmark for the reader.	Universally cited and **no primary source exists**. Traced through dozens of restatements to nothing. It is folklore.
Every fractional-executive adoption figure "1 in 4 US companies", "245% growth"	The market is moving toward senior operators.	Every one is published by a marketplace or staffing firm selling that exact service. None states a sampling frame. **No independent body has measured it.**
Acquisition cost by industry a widely reproduced agency table	A benchmark at the reader's exact vertical.	Method is "client analytics accounts, sourced anonymously", with **no sample size at all**, and the publisher's own service appears as the cheapest channel in its own table.
"40 to 60% of digital budgets are wasted"	The central argument of this guide.	Published by a consultancy selling waste reduction, on work from 2015 to 2018. We used the ANA log-level study instead, which discloses its method.
"$84bn lost annually to ad fraud"	A dramatic figure for the leakage section.	A model with no published methodology, and **contradicted by measured data**: median invalid traffic is 0.5% of spend.
"95% of buyers are not in market"	Why acquisition cost is structurally high.	Not rejected as false. It is a **heuristic** derived from purchase intervals, not a measured statistic, and quoting it as a research finding misrepresents what it is.

Note: One further figure, an agency's own funnel showing software crediting 78% of conversions to a channel that 12% of customers named unprompted, was the cleanest illustration of the attribution problem found anywhere. It is a single company's data, published by a firm that sells the thesis, and so it appears here rather than in the argument
Source: gRO source review, August 2026

ABOUT THE OPERATOR

Ro Maldonado

Founder, Growth Revenue Optimization

Fifteen years inside funnels, across B2B SaaS, fintech, wealth management, retail banking, and consumer brands. Award-winning across SEM, paid social, and integrated campaign work.

BEFORE GRO · PRIOR ROLES

$107M influenced pipeline at **27.57x** ROAS · **$2B+** sourced opportunities · demand gen owned for **$30B+** organisations. Prior-role credentials, not results my firm produced.

Where the model came from

I did not set out to start an agency. I set out to stop repeating myself.

In prior roles I was running **five products at once**, each with its own budget, its own funnel, and its own P&L, for Fortune 500 clientele on one side and early-stage startups on the other. Categories that had nothing in common. Buyers who wanted different things. Budgets that ranged from scrappy to millions.

The work that moved any of them was the same work every time. Find where the money was leaking. Fix the offer before touching the audience. Rebuild the measurement so the number meant something. Then run the experiment and read it honestly. What changed between products was the vocabulary. The method did not change at all.

Once that was clear, the question stopped being how many people it takes to run a growth function. It became how much of that function is **judgment** and how much is **production**. Judgment has to sit with someone who has been wrong before and remembers why. Production is everything downstream of a decision, and production is the part that finally scales without a hiring plan behind it.

gRO is that split, made into a company. One senior operator owns the judgment across the whole function, an agent fleet handles the production underneath, and the work happens inside your accounts against the number you already read.

$9,500 a month, all in.

Ad spend flows through your accounts. No markup, no media commission, the retainer is the only fee. It opens with a 90-day pilot. Four clients at a time.

gRO currently runs the growth function for five companies in five unrelated categories: a licensed real-estate brokerage, a career-tech platform, a workplace-safety practice, a venture studio, and a salsa brand heading to retail shelf. In each one the mandate is the whole function rather than a channel, and in each one it is the work described in this guide.

Figure 11: The mandate is the entire growth function, and it holds across five categories that share nothing

Current gRO engagements, the mandate in each, and what it has produced

ENGAGEMENT	THE MANDATE	WHAT IT HAS PRODUCED
RentList919 Licensed NC brokerage · live	Build renter acquisition from nothing and run it, end to end.	**$5.70 cost per lead at a 4.78% click-through rate**, verified against the platform's reporting API. Live and converting from day one.
OfferPath Career-tech SaaS, job-search · live	Own the entire top of funnel, including the sending infrastructure underneath it.	A brand system built from zero, a video engine shipping **40+ pieces in seven weeks**, an outbound programme operating at **six-figure list scale**, and sender reputation restored.
Vanguard Risk & Safety HSE advisory, South Texas · live	Fractional marketing leadership for a practice building its growth function from the ground up.	Brand identity, website, a **six-month marketing plan**, and SEO deployed, ahead of the first sales push rather than after it.
ScaleUp Ventures Multi-brand venture studio · live	Stand up paid acquisition, and a delivery model it can resell to its own client book.	Go-to-market plan, a **sourced funnel-economics model**, a full audit of the advertising estate, and a white-label delivery model. The audit caught a **platform change** that had removed the audience the plan relied on, before spend rather than after.
Tio Pelon's Salsa brand, retail shelf	Answer a positioning question with data instead of taste, before a retail launch.	**47 competitor labels coded across 14 variables** on a placement grid, and the flagship label rebuilt on the result. Four design corrections the data forced.

Notes: Named with each client's permission. Only the first row states an acquisition-cost outcome; every other cell is an output count, a deliverable, or a process finding, and none of it forecasts what any future engagement would produce
Source: gRO engagement records

The one acquisition number we put our name to

Exactly one figure above is a cost per acquisition, and it is ours: **$5.70**, on 71 leads and $404.35 of spend, on the RentList919 engine gRO built and still runs. It carries its denominator, its period and its source, which is all this guide asks of any number. Where we have no verified outcome we printed the work instead and claimed nothing.

Notes and sources

Every figure in this guide resolves to an entry below. Each was verified against the named primary source. Where a publisher reports a metric on a sub-sample, the sub-sample size is given rather than the headline participant count, because the two are rarely the same.

1 Gartner 2026 CMO Spend Survey, reported by *Marketing Dive*, June 2026. 401 CMOs across North America, the UK and Europe, surveyed January to March 2026. The majority of respondents are at companies above $1B in annual revenue. marketingdive.com

2 The CMO Survey, *Highlights and Insights Report*, February 2026. Duke Fuqua, Deloitte and the American Marketing Association. 2,111 US for-profit marketing leaders invited, 308 responded (14.6% response rate), 97% VP or above, fielded 7 to 29 January 2026. cmosurvey.org

3 Benchmarkit, *2025 B2B SaaS Performance Metrics Benchmarks*, May 2025. 583 participating companies in total, of which 72.1% are below $20M ARR. Sample size varies by metric and is stated on each exhibit. benchmarkit.ai

4 Nielsen, *2024 Annual Marketing Report*, 2024. Approximately 2,000 global marketers managing budgets of $1M or more, surveyed December 2023. nielsen.com

5 SaaS Capital, *2026 Spending Benchmarks for Private B2B SaaS Companies*, 2026. Fifteenth annual survey, more than 1,000 private B2B SaaS companies, completed March 2026. Denominator is ARR, not GAAP revenue. saas-capital.com

6 Aleph and Benchmarkit, *2026 SaaS and AI Performance Benchmarks*, June 2026. 342 companies, of which 198 reported CAC payback. getaleph.com

7 Meta Platforms, *Fourth Quarter and Full Year 2025 Results*, 2026. A filed financial disclosure covering all Meta advertisers, not a sample.

8 Tinuiti, *Digital Ads Benchmark Report Q1 2025*, 2025. Anonymised data from advertising programmes under Tinuiti management, representing more than $4B in annual spend. One agency's book of business, weighted toward retail and commerce. tinuiti.com

9 Association of National Advertisers, *Programmatic Media Supply Chain Transparency Study: Complete Report*, December 2023. 21 marketers and 12 supply-chain companies, $123M of spend and 35.5 billion impressions analysed at log level between September 2022 and January 2023, with 35 supporting interviews conducted by Kroll. Agency fees were explicitly outside the study's scope. ana.net

10 Association of National Advertisers, TAG TrustNet and Fiducia, *Programmatic Media Quality Benchmark*, Q1 2025 and Q1 2026 editions. Q1 2025: 39 marketers, $242M of spend, 41.9 billion impressions, November 2024 to March 2025. Q1 2026: 86 participating marketers, 66 active contributors. ana.net

---

Figures attributed to **gRO analysis**, **gRO illustration** or **gRO method** are the firm's own work and are labelled as such at the point of use. The worked example on page 11 is arithmetic on stated inputs and does not describe any client. Claims about gRO's own results, and about the founder's record before gRO, are separated throughout and carry their own labels.