Metadata

FOR THE CMO AND VP OF MARKETING

What *$57.6M* of B2B ads says about your budget

Six numbers that change where the money goes next year — and what each one still cannot tell you.

$57.6M
of 2025 B2B ad spend

153
advertisers

211,000
leads, matched to their CRM

- 3 in 4 paid more than $202
- 62% cannot trace a dollar
- $35,288 vs $130,468 a customer
- $12.7M bought 66 leads
- 25% / 7% / 16% budget split
- $60,896 vs $69,705 a customer

Prepared by Metadata.io • Gil Allouche, Founder & CEO

metadata.io/benchmark-report-2026

01 • THE NUMBER IN EVERY DECK

Three out of four advertisers paid more than *$202*.

Count each advertiser once instead of counting dollars, and the one in the middle paid $376 for a LinkedIn lead. The figure everyone quotes is a spend average, so the biggest spenders drag it down.

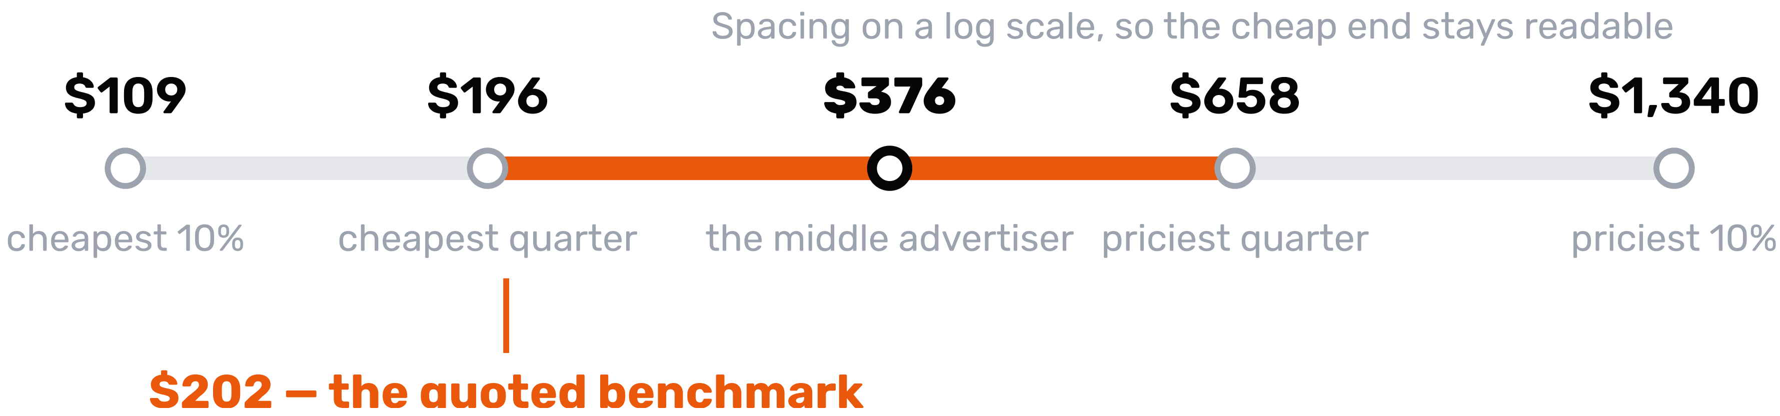

The catch. Across 88 LinkedIn advertisers. Both numbers are true and they answer different questions: $202 is what the market paid per dollar, $376 is what a company like yours paid.

WHAT TO DO ABOUT IT

If your cost per lead sits above $202, most of this market is there with you. Bring this scale to the budget conversation before someone brings a blog post instead.

02 • THE MEASUREMENT PROBLEM

62% of advertisers cannot see what their ads returned.

Of 105 advertisers running lead generation, 65 had no closed-won revenue their CRM could trace to an ad twelve months on. Not weak returns — no visible returns at all.

Share of advertisers running lead generation in 2025

The catch. Among the 39 who can trace it, attributed payback was 0.66x inside twelve months. Across all 105 it reads 0.54x — but that number is mostly measuring who has the plumbing, not whose ads worked. B2B deals also close later than a twelve-month window.

WHAT TO DO ABOUT IT

Before you judge a channel on payback, check whether closed-won revenue can reach it at all. If it cannot, you are not measuring performance — you are measuring your own attribution coverage.

03 • WHO YOU CHASE

A mid-market customer cost *$130,468*.

The same money bought a customer for $35,288 when the target was a 51–200-employee company on LinkedIn. One blended cost-per-customer number hides a gap this size completely.

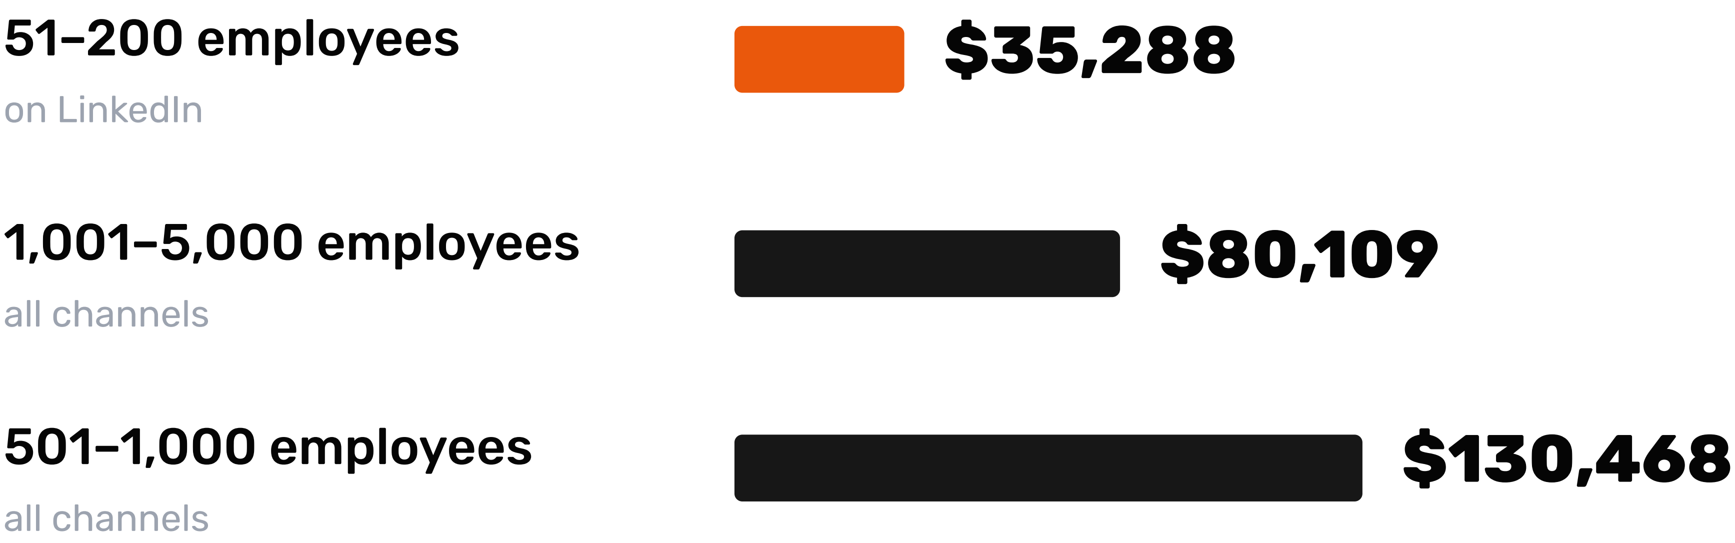

The catch. Different advertisers chase different sizes, so part of this is who they are rather than who they targeted. The gap is far too wide to be only that.

WHAT TO DO ABOUT IT

Score small-company campaigns on payback. Score enterprise on coverage instead. Judging both on the same number is how a working enterprise programme gets cut.

04 • THE WEAKEST LINE IN THE PLAN

$12.7M bought clicks and produced 66 leads.

That is 22% of every dollar here, across 112 advertisers. 61 of the 153 put more than 30% of their budget into it.

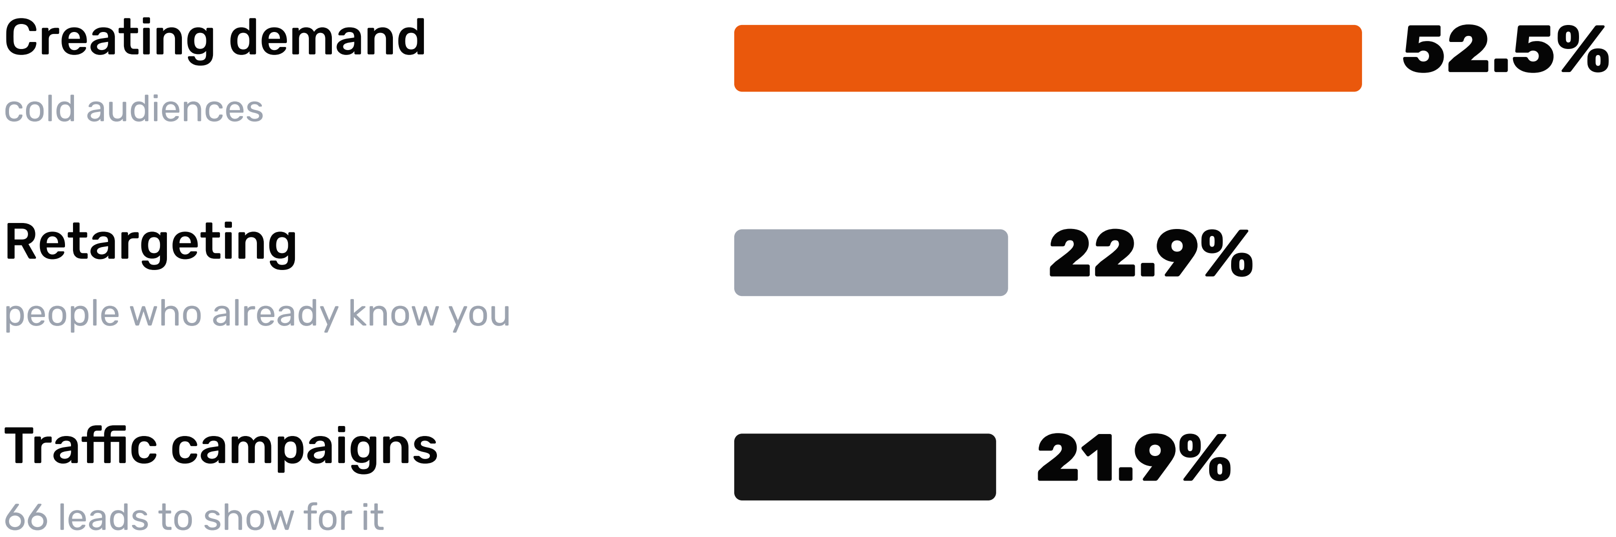

Share of every dollar in this dataset

The catch. A traffic campaign has no form on it, so lead reporting has nothing to count. This is a broken measurement, not proof the money vanished — and whatever brand value it built, we did not measure that either.

WHAT TO DO ABOUT IT

Every traffic campaign gets a written purpose, a number it has to hit, its own budget line, and no place in your cost-per-lead report. If nobody can name what one is for, stop paying for it.

The middle advertiser puts *25%* into creating demand.

And 7% into retargeting, 16% into traffic campaigns. The spread is enormous: on every line the boldest quarter spends many times what the quietest quarter does.

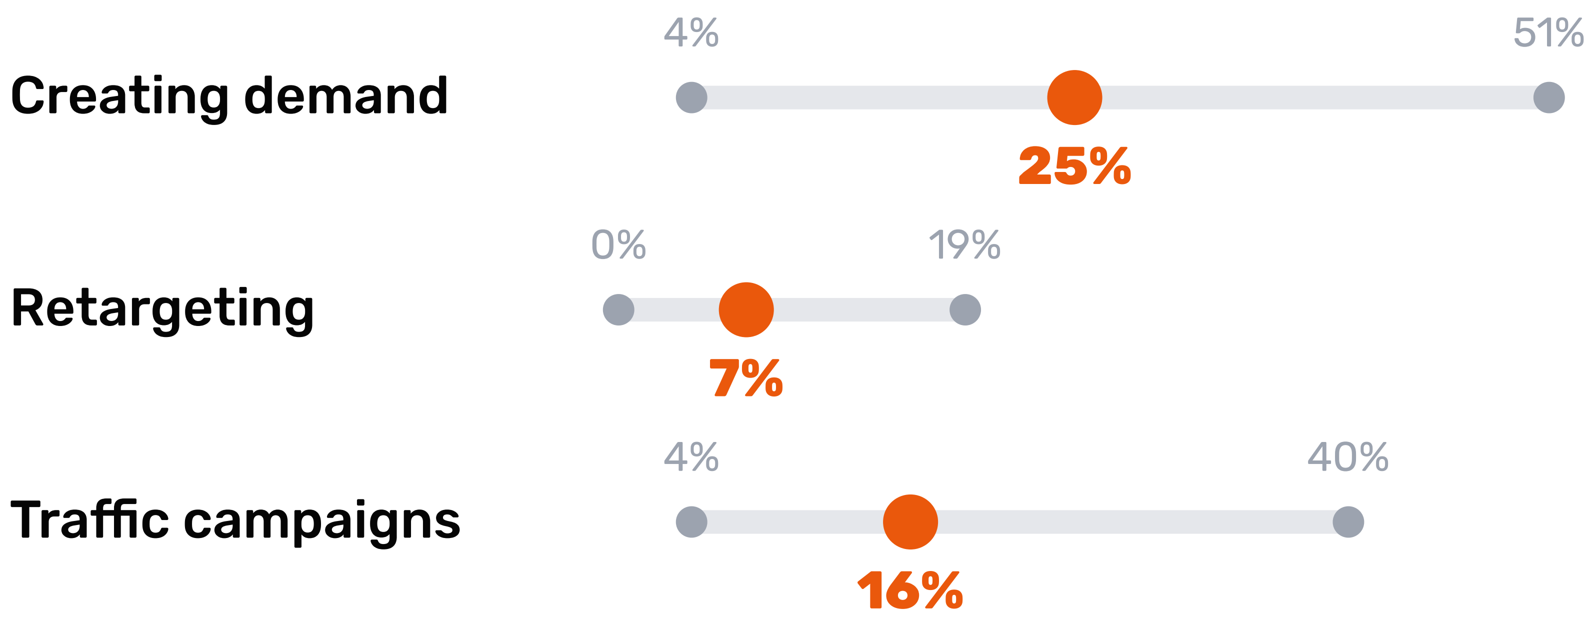

Grey ends: the quietest and the boldest quarter. Orange: the middle advertiser.

The catch. Across 109 advertisers who spent more than $10,000. This is what the market did, not what worked: the same advertisers only got 0.56x back.

WHAT TO DO ABOUT IT

Put your own three percentages next to these before the planning meeting. Then argue about what each line returns, not about matching a number the market happened to land on.

Higher-cost leads produced the cheaper customers.

Advertiser-built audiences cost $217 a lead against $181 for the platforms' own audiences — then produced 3.56 customers per thousand leads against 2.6.

The catch. Platform-native audiences still returned more revenue per dollar, because the deals it brought were bigger. We cannot separate that from who those advertisers were selling to, so we claim the cheaper customer and not the bigger deal.

> **WHAT TO DO ABOUT IT**
>
> **If your audience budget is decided on cost per lead, it is being decided on the wrong number. Put cost per customer on that report.**

Five questions for your next team meeting

None of these needs new software. All five are answerable from what you already have.

01 What share of paid budget is in traffic campaigns, and what is each one meant to produce?

Pull: Spend on traffic-objective campaigns, the stated goal of each, and the metric it is judged on.

Good: Every traffic campaign has a written purpose and a success metric that is not cost per lead.

Bad: Traffic spend is judged on leads, or has no stated goal. The middle advertiser here put 16% of budget there.

02 What does a customer cost us by the size of company we target?

Pull: Media spend and closed-won customers split by employee band, then cost per customer for each.

Good: Cost per customer is reported per segment, and the segments have different targets.

Bad: Enterprise and SMB are judged against one blended number. Here they ranged from $35,288 to $130,468.

03 Where does our LinkedIn cost per lead sit against what other advertisers paid?

Pull: LinkedIn lead-generation spend divided by leads, for the same twelve months.

Good: Between $196 and $658 puts you in the middle half of this market.

Bad: Treating $202 as the target. It is a spend-weighted average; the middle advertiser paid $376.

And then these

04 Can closed-won revenue in our CRM be traced back to a campaign at all?

Pull: Closed-won opportunities with a campaign or channel attached, as a share of all closed-won.

Good: Campaign, lead, opportunity and closed-won are joined, with open pipeline reported separately.

Bad: Nothing connects. 62% of the advertisers here were in that position.

05 Would our conclusions survive losing our single biggest customer?

Pull: Recompute cost per customer and payback with the largest closed-won deal removed.

Good: The answer barely moves, and you report the median next to the average.

Bad: One deal changes the story. It is why every figure in this report is capped so no single advertiser can swing it.

All of it, filterable by your company size, industry and channel:
metadata.io/benchmark-report-2026

JUMP TO ANY PAGE

What is in here

Every number here is filterable by company size, industry and channel at **metadata.io/benchmark-report-2026**

HOW WE COUNTED

Five rules we held ourselves to

We only count deals your CRM counts.

Every revenue figure is matched to closed-won in the advertiser's own CRM. What the ad platform claims it drove is never used.

Twelve months, and we say so.

Deals still open when we froze the data count as zero. That makes every revenue number here a floor, not a forecast.

When several ads touch one deal, they split the credit.

Equally. So the numbers still add up when you roll them together.

If a number rests on too few advertisers, we do not print it.

It needs at least 5 advertisers and $50,000 of spend behind it, and no single advertiser can account for more than half. Anything that fails is left out and named, never estimated.

When we get something wrong, we say so in public.

Twelve claims have been killed or corrected so far, each with its reason on the site.

The full method, every definition and every correction:
metadata.io/benchmark/methodology

READ THIS BEFORE YOU QUOTE US

What this data cannot tell you

We cannot tell you the ads caused the sale.

Nobody ran a holdout. This is ad-attributed closed-won revenue, which is not the same thing as revenue the ad created.

We cannot tell you what you should spend.

The budget splits describe what advertisers did in 2025. Plenty of them were wrong.

We cannot show you a trend.

This is one year of data. Ask us again next year.

We cannot tell you your profit.

Costs here are media only. No agency fees, no salaries, no production.

We cannot speak for advertisers who are not our customers.

Everyone in this data runs ads through Metadata and let us match them to their own CRM.

See how your size, industry and channel compare

metadata.io/benchmark-report-2026

Free, filterable, no email required.