B2B Marketing KPIs

Fifteen marketing metrics with named-source benchmarks, including the industrial and channel ones nobody publishes

By Andre Magrini, fractional Chief Revenue Officer, Chicago. Every benchmark in this document carries a named source. Where no credible public benchmark exists, it says so instead of inventing one.

andremagrini.com

Almost every article on marketing KPIs was written for a software company. It assumes a buyer who finds you through search, downloads something, gets scored, gets nurtured by email, and buys in ninety days. That buyer exists. He is not your buyer.

If you sell capital equipment through dealers, the standard KPI stack was built for different physics. The cycle is nine to eighteen months. The buyer trusts his dealer more than your website. Half your demand never touches a form. And your most valuable marketing asset is a relationship between a regional manager and a dealer principal who has sold your line for twenty years.

Fifteen KPIs follow, with named sources where public benchmarks exist and an honest "no reliable public benchmark" where they do not. Each includes a section on how it lies to you, which is what most KPI content omits and the part that saves money.

Part 1: The two KPIs that start every fight

1. Marketing Qualified Lead (MQL)

What it is. A contact marketing has decided is worth a salesperson's time.

The formula. There is none, and that is the problem. The closest thing: a contact whose lead score crosses a threshold your team chose.

Benchmark. Forrester analyst Terry Flaherty, in "The Revenue Process Alignment Series, Part 1: The End Of MQLs" (April 2022), reported that conversion from inquiry to closed-won in a lead-centric process is typically under 1%, and that over 80% of B2B buying decisions are made by groups of more than three people. First Page Sage reports 13% MQL-to-SQL for B2B SaaS from client data collected 2019 to 2025, and that number does not describe industrial or channel businesses. No credible public MQL-to-SQL benchmark exists for agricultural equipment, capital goods, or dealer-network manufacturing. Anyone quoting one is quoting a SaaS figure with the label changed.

How this KPI lies to you. An MQL is not a measurement, it is a definition, and a definition is controlled by whoever is measured on it. Pay a team on MQL volume and you have asked them to widen a definition, not to create demand. They will: drop the threshold ten points, count a webinar registration, count a badge scan. Volume rises 40% in a quarter, revenue does not move, and every individual decision was defensible. In a dealer network there is a second failure: the model assumes the qualified person is the buyer, when frequently the real buyer is already talking to your dealer, where your automation cannot see him.

What to do with it. Keep the MQL as a routing rule and remove it from the incentive plan and the board deck. In a dealer model the honest replacement is opportunities accepted by a dealer, the first moment someone with revenue accountability agreed the lead was real.

2. Sales Qualified Lead (SQL)

What it is. A lead sales has inspected and agreed to work.

The formula. MQL-to-SQL conversion = SQLs accepted / MQLs delivered.

Benchmark. No reliable public benchmark exists for dealer-distributed capital equipment. The First Page Sage SaaS figure above is the most transparently sourced public number, and it does not transfer. Use your own trailing twelve months.

How this KPI lies to you. It looks like a quality measurement. It is a negotiation outcome. If sales has capacity, acceptance rises. If reps are in the field during harvest, it falls. If someone is behind on coverage, it rises. So when the rate moves from 18% to 26%, the first question is not what marketing did better, it is what changed about who accepts. The deeper point: the definition of SQL is political, not technical. It is a boundary between two budgets, so wherever you draw it one department gets credit and one gets blame, and the line drifts toward whoever has more organizational power.

What to do with it. One definition, one document, signed by both leaders, with examples and counterexamples, then frozen for four quarters. A mediocre definition held constant produces usable trend data. An excellent one that changes quarterly produces nothing.

Part 2: Contribution and cost

3. Marketing Contribution to Pipeline

What it is. The share of pipeline marketing helped create, split into sourced (marketing was first touch) and influenced (marketing touched the deal at any point).

The formula. Marketing-sourced % = pipeline from opportunities whose first recorded touch was a marketing activity / total new pipeline.

Benchmark. What is documented is Forrester's position, not a benchmark. In March 2022 Forrester principal analyst Brett Kahnke published "Marketing Sourcing: A Metric Built to Get a CMO Fired," arguing the metric is structurally broken: demand marketing is a third or less of a typical marketing program budget, and roughly three-quarters of B2B revenue comes from existing clients, which pushes the sourced number down over time regardless of performance. No reliable public benchmark exists for marketing-sourced pipeline in industrial or channel businesses. The 30% to 50% figures circulating online are attributed to Forrester and Gartner studies that are difficult or impossible to locate.

How this KPI lies to you. Sourced pipeline measures instrumentation, not contribution. Gartner's B2B buying journey research finds buyers spend about 17% of purchase time with all suppliers combined, and as little as 5% to 6% with any single one. 6sense's 2025 B2B Buyer Experience Report found first seller contact happens at 61% of the journey, and that roughly 95% of winning vendors were already on the Day One shortlist. Bain, in research with Google covering 1,208 US buyers across categories that included industrial equipment, found 80% to 90% have a vendor set in mind before any research begins. The decisive work happened before anything measurable occurred. And if the deal enters through the dealer, marketing sources 0% of it by definition, even when the reason that farmer walked in was eight years of your brand at every field day in the county.

What to do with it. Report sourced and influenced together, and label sourced honestly as "pipeline where marketing was the first tracked touch." Add one self-reported question to every quote request and dealer intake: how did you first hear about us. Once a year, cut brand spend in one of two matched territories and compare inquiry volume twelve months later.

4. Customer Acquisition Cost by Channel

What it is. What one new customer costs through a specific channel, including people cost, not just media.

The formula. (Media + agency and production + fully loaded internal labor for that channel) / new customers acquired through it.

Benchmark. First Page Sage's B2B CAC by industry report, from client data collected January 2022 to August 2025, gives Manufacturing at $662 organic, $905 inorganic, $723 blended; Agtech at $712; Industrial at $542. Two caveats. The blended figure is weighted 75% organic to reflect that firm's own service mix, not the market's. More importantly, the dataset excludes email, in-person events, direct mail, and outdoor advertising for insufficient data, which excludes the channels generating most industrial demand.

How this KPI lies to you. You cannot compute cost per customer by channel without first deciding which channel gets credit, so every channel CAC is an attribution model wearing a finance costume. It lies again through timing: spend is booked in the quarter it happens, customers arrive nine to eighteen months later, so the metric looks terrible when you increase investment and excellent when you cut it. And dealer-sourced customers show near-zero marketing CAC because no marketing cost was tagged to them, so the dealer channel looks free. It is not. It costs margin points, co-op accruals, territory exclusivity, and field support headcount.

What to do with it. Compute CAC on a cohort lag matching your real cycle, and load dealer margin, co-op, and field support into a dealer-channel CAC so channels compare on the same basis. A $9,000 CAC that opens a dealer relationship worth twenty years of parts revenue is cheaper than a $900 CAC that sells one machine.

5. Blended CAC versus Paid CAC

What it is. Blended divides all acquisition cost by all new customers, including those who would have arrived anyway. Paid divides paid media cost by only paid-attributed customers.

The formula.

- Blended CAC = total sales and marketing cost / all new customers
- Paid CAC = paid media and associated cost / customers attributed to paid channels

Benchmark. The most useful public anchor is a spend figure, not a CAC figure. Gartner's 2025 CMO Spend Survey (12 May 2025, n=402 marketing leaders across North America, the UK and Europe, most above $1B revenue) found marketing budgets flat at 7.7% of company revenue, with paid media at 30.6% of marketing budget, or 2.4% of company revenue. Half of CMOs reported 6% or less.

How this KPI lies to you. They lie in opposite directions, which is why running only one is dangerous. Blended CAC flatters paid media, because it buries paid spend in a denominator containing every dealer referral, repeat buyer, and twenty years of reputation. Double paid spend, generate nothing incremental, and blended barely moves. Paid CAC flatters organic, because it assumes every paid-tagged customer was caused by the channel, when many people clicking a branded search ad would have found you anyway. That is not acquisition, it is paying a toll on demand you already created.

What to do with it. Report both and read the gap as the diagnostic. Then run an annual geographic holdout. Dealer territories are already clean geographic boundaries with their own sales data, which makes incrementality testing far easier for you than for a software company. Use that advantage.

6. Cost per Qualified Opportunity

What it is. What it costs to produce one opportunity a salesperson accepted. It sits between cost per lead (nearly useless) and CAC (accurate but too late to steer with).

The formula. Total marketing cost in the period / qualified opportunities created in the period.

Benchmark. No reliable public benchmark exists, in any industry. A benchmark requires a shared definition, and this denominator is defined differently by every company reporting it. Published figures are single-vendor client averages or unattributed. Build a baseline over four quarters and compare only to yourself.

How this KPI lies to you. It lies exactly as MQL does, but with a dollar sign attached, which makes it look like finance: because the denominator is a judgment call, the cost improves the moment you loosen the criteria. It lies again through lag. Cut the budget and cost per opportunity improves, because the pipeline still holds opportunities last year's spend created. That improvement is a mirage, and it usually arrives one or two quarters before the real problem does.

What to do with it. Use it as a directional early warning, quarterly, never monthly, segmented by product line and region. Require that any reported improvement carries a statement that the opportunity definition did not change.

Part 3: The KPIs of motion

7. Pipeline Velocity

What it is. How many dollars of expected revenue move through the pipeline per day. The best single diagnostic in a long-cycle business, because it forces four levers into one visible number.

The formula. (Qualified opportunities x average deal value x win rate) / average sales cycle length in days.

Benchmark. No reliable public benchmark exists, and one could not, because the output is denominated in your deal sizes and your cycle. The nearest external anchor is cycle length: 6sense's 2025 B2B Buyer Experience Report puts the average B2B buying cycle at 10.1 months, down from 11.3 in 2024, with buying groups of roughly ten or more people. Dealer-sold capital equipment typically runs longer, and no public benchmark exists for the agricultural equipment cycle specifically.

How this KPI lies to you. Four variables, one output, so four different stories produce the same number. Velocity rising on improved win rate is a healthy business. Velocity rising because deal value grew while opportunity count collapsed is a business about to run out of pipeline. The most common self-deception is the denominator: most CRMs measure cycle length only on closed deals, so stalled deals are excluded and velocity is overstated. In seasonal capital equipment the distortion is large, because deals that miss the buying window do not die, they sleep, and sleeping deals never enter the average.

What to do with it. Never report velocity alone; show all four inputs beside it. Measure cycle length on all opportunities including aged and lost ones, and compare to the same quarter last year. In most industrial businesses the cheapest lever is cycle length, and the fix is usually a specification or financing bottleneck rather than anything marketing controls.

8. Stage Conversion Rate

What it is. The percentage of opportunities advancing from one pipeline stage to the next.

The formula. Opportunities advancing to stage N+1 / opportunities that entered stage N, measured on a closed cohort.

Benchmark. The one durable public figure is Forrester's (Flaherty, April 2022): inquiry to closed-won in a lead-centric process is typically under 1%. Stage-by-stage figures published online vary so widely

(MQL-to-SQL quoted anywhere from 13% to 40% in the same year) that variance exceeds signal. No reliable public stage-conversion benchmark exists for industrial or dealer-channel businesses.

How this KPI lies to you. It is measured on a cohort that has not finished. If a stage takes five months to clear and you measure a cohort that entered three months ago, conversion looks catastrophic because most of the cohort is still sitting there, and teams then fix an artifact of measuring too early. The costliest lie is the reverse of the obvious one: a high conversion rate usually means the entry criteria are too strict. If 85% of your proposals close, you are not writing great proposals, you are refusing to quote deals you could have won.

What to do with it. Measure only on fully closed cohorts. Track time in stage alongside conversion, because a stage converting at 70% but sitting seven months is your actual bottleneck. Investigate high rates with the same suspicion as low ones.

9. Attribution

What it is. The assignment of revenue credit to marketing touchpoints. Last-click gives all credit to the final touch; multi-touch splits it by a rule.

The formula. There is no correct formula, only chosen conventions. Last-click assigns 100% to the last recorded touch. W-shaped assigns 30/30/30 to first touch, lead creation, and opportunity creation. Every weighting is a convention, not a finding.

Benchmark. No reliable public benchmark exists, and the absence is itself the finding. No credible study establishes attribution accuracy, because doing so would require knowing the true causal contribution, which is the thing attribution was invented to estimate. What is well sourced is how much of the journey is invisible: Gartner's 17% of purchase time with all suppliers and 5% to 6% with any one; 6sense's 2025 finding that first seller contact happens at 61% of the journey; Bain's finding that 80% to 90% of buyers have a vendor set in mind before research begins.

How this KPI lies to you. Last-click attribution in a nine-month cycle does not slightly misallocate credit. It systematically inverts your budget.

Late touches happen when a buyer already knows he wants your product: branded search, a pricing page, a dealer locator, a quote form. They are cheap, measurable, and nearest the conversion, so last-click loads them with credit. The early touches that put you on the Day One list nine months earlier (a field day, a magazine spread, a dealer's recommendation, your machine working in a neighbour's field) are expensive, untrackable, and get zero. So the model tells you to defund what creates demand and refund what harvests it. Follow it for two years and you own an efficient harvesting operation with nothing left to harvest, and the metric looks excellent throughout. Last-click failure is invisible in the last-click report. Multi-touch is better but not different in kind: it sees only instrumented touches, so it cannot see the conversation in the dealer's parts department.

What to do with it. Demote attribution to a directional report, then build three things that actually inform budget: self-reported attribution on every inbound and dealer intake, geographic incrementality tests using dealer territories as natural cells, and an annual brand tracking survey with identical questions year over year. The third is the only instrument that sees the roughly 95% of the market not currently buying.

10. Share of Voice and Excess Share of Voice

What it is. Share of voice is your share of total category advertising or conversation. Excess share of voice (ESOV) is share of voice minus share of market, and it is the version that predicts anything.

The formula.

- Share of voice = your category advertising spend / total category advertising spend
- ESOV = share of voice minus share of market, in percentage points

Benchmark. One of the best-evidenced relationships in marketing. Les Binet and Peter Field, analysing the IPA databank, established that brands whose share of voice exceeds share of market tend to grow and those below tend to decline, at roughly 0.5 percentage points of annual market share growth per 10 points of sustained excess share of voice in consumer categories. In their work for the LinkedIn B2B Institute ("The 5 Principles of Growth in B2B Marketing", 2019) they found the relationship holds in B2B, with effectiveness maximised at approximately 46% of budget to brand building and 54% to activation. The supporting logic comes from the Ehrenberg-Bass Institute: John Dawes' work for the LinkedIn B2B Institute established the 95:5 rule, that at any moment roughly 95% of business buyers are not in the market, because firms replace most supplier categories only every few years.

How this KPI lies to you. It lies through the denominator. Measured across "agriculture" you compete with Deere's global spend and your share rounds to zero. Measured across "grain handling equipment in the US Midwest" you might hold 15%. Whoever picks the category boundary picks the answer. It lies again by ignoring the channels that carry your voice in a channel business: share of voice tools count paid media and online mentions, not twelve dealer showrooms with your banner in the window, the field days where your machine ran, or the trade magazine your buyers actually read.

What to do with it. Define the category as narrowly as you honestly compete and hold that definition constant. Include co-op-funded dealer advertising in the numerator, because your buyer does not distinguish your ad from your dealer's ad carrying your logo. If share of market sits well above share of voice, you are harvesting equity someone built earlier, and the bill arrives later.

Part 4: The KPIs that exist only in industrial and channel marketing

11. Trade Show Return

What it is. Revenue and pipeline generated by exhibiting, against the fully loaded cost of exhibiting.

The formula. (Pipeline attributable to the show x historical win rate x gross margin) / total show cost, where cost includes space, build, freight, drayage, travel, lodging, staff time at fully loaded rates, and pre-show and post-show campaign cost.

Benchmark. The exhibition industry circulates dramatic figures that do not survive checking. The widely quoted "$20.98 returned for every dollar" and "$142 per trade show lead" are attributed constantly to CEIR (the Center for Exhibition Industry Research) with no dated, locatable primary report behind them. Do not put those in front of a CFO. What is verifiable from CEIR: its "Role and Value of Face-to-Face Interaction" study series, in the report "Profiles of Attendees and Exhibitors" (December 2012, based on more than 9,000 attendees and 800 exhibitors), found 94% of attendees possess net buying influence, meaning they specify, recommend, determine need, influence, or decide. That is a real finding from a named organisation with a stated sample, and it is now over a decade old. There is no current reliable public benchmark for trade show ROI or cost per show lead.

How this KPI lies to you. It is usually calculated as show cost divided by badge scans, which measures foot traffic, and many scans are existing customers and open opportunities, which inflates the show and deflates every other channel. The largest lie is the denominator: most calculations use the booth invoice and miss freight, drayage, the hotel block, and the fully loaded cost of eight people off territory for four days, which in capital equipment can turn a "3x return" into a loss. And most show value is destroyed afterward. The Harvard Business Review study by Oldroyd, McElheran and Elkington ("The Short Life of Online Sales Leads", March 2011) found conversion probability roughly 60 times higher when firms respond within an hour, yet show leads routinely sit two weeks in a spreadsheet before reaching a dealer.

What to do with it. Load the full cost including staff days. Segment scans into existing customers, open opportunities, and genuinely new accounts, and calculate return only on the third. Set a follow-up SLA in hours and report compliance as its own KPI. Measure on a twelve-month tail, not the sixty days after the doors close.

12. Field Day Return

What it is. The return on demonstrating equipment in working conditions on a customer's or dealer's site. In agricultural and heavy equipment markets this is often the highest-conviction marketing activity available, and among the most expensive per attendee.

The formula. (Quotes issued within 180 days to attendees x win rate x gross margin) / total event cost including machine transport, machine downtime, host compensation, staff time, and hospitality.

Benchmark. No reliable public benchmark exists. No industry body publishes demonstration-day benchmarks; CEIR's research covers organised exhibitions, not on-farm or on-site demonstrations. Every field day ROI figure in circulation is one company's internal number or an agency's claim.

How this KPI lies to you. Attendance is organised by the host dealer, so the room is full of that dealer's existing customers, and you spent the day marketing to people who already buy from you. Attendance count is close to worthless; attendance of accounts not currently buying from you is the metric. The payoff timing lies too: an August field day produces quotes the following February and March, so any window shorter than a full season underreports it. And the largest effect runs through dealer commitment, not end-customer conversion. A dealer who hosted your demo has publicly attached his reputation to your line, and it shows in his ordering for two years. That value never appears in an event ROI calculation.

What to do with it. Track new-account attendance (no purchase from you in 24 months) and quote conversion within 180 days, evaluated on a full season. Then track the host dealer's order volume for eight quarters against a matched dealer who did not host.

13. Dealer Co-op Utilization and Return

What it is. Co-op is the fund you accrue for dealers, usually a percentage of purchases, to spend on local marketing of your line. Utilization is how much of it actually gets claimed and spent.

The formula.

- Co-op utilization = co-op dollars claimed and approved / co-op dollars accrued
- Co-op return = incremental dealer revenue on your line during the campaign window / co-op dollars spent

Benchmark. The circulating figure is sourceable but old. The Local Search Association white paper "Breaking Through the Co-op Clutter," reported by MediaPost (Laurie Sullivan, 7 November 2016), estimated that of $36 billion to $70 billion in annual co-op advertising dollars available from US manufacturers, roughly $14 billion to $35 billion goes unclaimed. That is an estimate, nearly a decade old, spanning all of US retail and manufacturing. No current reliable public benchmark exists for co-op utilization in agricultural or industrial equipment.

How this KPI lies to you. Utilization is treated as marketing performance when it is almost entirely a measure of programme friction. Low utilization rarely means dealers do not want the money. It means your claim process requires pre-approval, proof of performance, tear sheets, and a portal login, and a dealer principal with eleven employees decided it was not worth the paperwork. You are measuring your own bureaucracy and calling it dealer engagement. It lies upward too: co-op is easiest to spend on what is easiest to claim, typically price-led local advertising, so a fully utilised budget often means you financed a season of discounting for your own product. And because co-op accrues on what a dealer already bought, the fund rewards past performance and does almost nothing to develop the dealers you need to grow.

What to do with it. Report utilization by dealer tier, not in aggregate. Track claim cycle time (days from submission to reimbursement) as the leading indicator, because that drives utilization more than dealer enthusiasm does. Cap the share of claims going to price-led advertising, and measure return against a matched dealer control.

14. Specification Generation (Spec-in Rate)

What it is. The rate at which your product is named in a specification, drawing, tender, or approved-vendor list before the purchase decision. In capital projects and OEM supply, getting specified is frequently the moment the deal is decided and everything after is paperwork.

The formula.

- Spec-in rate = projects where you are named basis of design or approved manufacturer / identified projects in your addressable segment
- Spec-to-order conversion = orders won on specified projects / projects where you were specified
- Spec displacement rate = projects where you were specified but a substitute was approved and won

Benchmark. No reliable public benchmark exists. Dodge Construction Network and ConstructConnect maintain the project and specification databases that make specification tracking possible in construction-adjacent markets and are the credible named sources for the underlying project data, but neither publishes a validated public conversion benchmark. Figures such as "specifications convert to orders at 30% to 60%" appear on vendor blogs with no study behind them. Do not cite them.

How this KPI lies to you. Most dangerously through substitution. Being specified feels like winning and the report says you won, but most tender structures permit an approved equal. If your specification is written loosely enough that three competitors qualify, you did the engineering and someone else took the order on price. A company can hold a rising spec-in rate and a falling win rate at the same time, and the dashboard looks excellent throughout. Second, the denominator is almost always wrong: you can only compute a rate against projects you know about, and in fragmented industrial markets that is the subset where you already have relationships. It is biased upward by exactly the amount of market you cannot see.

What to do with it. Track spec-in and spec-to-order together and treat the gap as the real metric. Classify every specification as exclusive (no equal permitted), basis of design (yours named, equals permitted), or listed. Only the first is genuinely defensible, and the mix tells you more than the total. Report on a three-year chart, because specifications written this year become orders in eighteen to thirty-six months.

15. Dealer Lead Follow-up Rate

What it is. The percentage of leads passed to a dealer that receive a documented contact attempt inside your agreed window. This is the channel equivalent of lead response time, and usually the largest single point of leakage between marketing spend and revenue.

The formula. Leads with a documented dealer contact attempt within the SLA window / total leads passed to that dealer.

Benchmark. No reliable public benchmark exists for dealer-network follow-up specifically. The relevant sourced evidence is on response speed generally: the Harvard Business Review study by James Oldroyd, Kristina McElheran and David Elkington (March 2011) found conversion probability roughly 60 times higher when a firm responds within an hour, and documented that most firms were far slower than they believed.

How this KPI lies to you. It lies by omission, because in most manufacturers it is not measured at all. Leads go into a portal and vanish from view. Marketing reports leads delivered, nobody reports leads worked, and the gap is frequently enormous. Where it is measured, it lies through self-report: a "contacted" flag with no timestamp measures a checkbox, not a response time.

What to do with it. Require a timestamped disposition in a shared system before a dealer receives the next lead, and publish follow-up rate by dealer across the network, because visible comparison changes behaviour faster than policy does. Measure this before you spend another dollar on leads: if your network works 40% of what you send, doubling volume is the most expensive possible way to fix a problem a process change would fix for free.

The vanity KPIs that cost money

A vanity metric is not simply a number that fails to correlate with revenue. Plenty of harmless numbers do that. A vanity metric is one that redirects real budget and real behaviour toward activity that produces nothing.

Metric	Why it costs money, not just attention
Impressions and reach	Priced per thousand, so the cheapest way to improve it is to buy the least relevant audience available. Optimising for reach in a market of 4,000 real buyers actively funds waste.
Website sessions	Rewards content that attracts students, competitors, and job seekers. Teams measured on traffic build a blog ranking for terms your buyers never search.
Followers and engagement rate	Where the buying population is a few thousand people, follower count is dominated by people who will never buy, and engagement rate improves as the audience gets less qualified.
Email open rate	Structurally unreliable since Apple Mail Privacy Protection began pre-loading images in 2021, registering opens no human performed. Close to noise now, and it still drives subject-line work.

Metric	Why it costs money, not just attention
Gross lead count and cost per lead	The most expensive vanity metric in industrial marketing. Cheap leads are cheap because they are unqualified. A team held to cost per lead buys the worst available traffic, hits target every quarter, and teaches dealers to stop trusting the leads.
Content volume	Measures throughput of an internal process, not market effect. Funds a production line whose output nobody has verified anyone reads.
Trade show badge scans	Justifies show budget on foot traffic, and encourages staff to scan indiscriminately rather than qualify, degrading the lead file and the follow-up rate at once.
MQL volume	A metric whose definition is controlled by the team incentivised on it will always improve, and the improvement means nothing.

The test for anything on your dashboard: if this number doubled next quarter and nothing else changed, would the business be better off? If you cannot answer yes with a straight face, it does not belong in a review meeting.

How marketing and sales stop fighting over the number

Almost every marketing-versus-sales argument in a channel business is the same argument in different clothes: two functions are measured on two definitions of one event, and both definitions are correct inside their own system. Marketing says it delivered 300 qualified leads. Sales says it received garbage. Both are reading accurate reports describing different things. This is not fixed with better feelings or a shared dashboard. It is fixed with three artifacts.

One definition of each stage, written down and frozen

For each pipeline stage, specify:

1. Entry criteria as observable facts, not judgments. "Buyer confirmed a budget range and a purchase window" is observable. "Buyer seems serious" is not.
2. Who may move a record into the stage, by role.
3. Two worked examples that qualify and two that do not. Counterexamples do more work than definitions.
4. What happens to a record that fails: returned to marketing, recycled, or closed. Every rejected lead needs a documented destination, or rejections vanish and the disagreement becomes unresolvable.

Then freeze it for four quarters. Imperfection matters far less than stability, because stability is what makes trend data mean anything. Review on a set date, never during a bad month, and never as the outcome of an argument.

A two-way SLA with numbers on both sides

Most lead SLAs are one-directional: marketing promises volume, sales promises nothing. That is not an agreement, it is a quota with extra steps.

Marketing commits to	Sales and dealers commit to
A defined number of qualified opportunities per period, by territory and product line	First documented contact attempt within a stated window (24 business hours is a realistic starting point in a dealer network; the HBR evidence argues for faster)
Every lead delivered with a minimum data set: source, self-reported origin, product interest, timeframe	A minimum number of contact attempts across at least two channels before disposition
Delivery through one agreed system, not email, not a spreadsheet	A disposition code recorded within a stated window, from a fixed list, with a reason for every rejection
A named owner for rejected-lead review	Participation in a monthly rejected-lead review

Add one consequence clause that makes it real: leads not dispositioned within the window return to the central pool and can be reassigned. In a dealer network that is the only enforcement mechanism that reliably changes behaviour, because it touches territory rather than pride.

One shared number that neither side can move alone

The structural cause of the fight is that marketing is measured on lead volume and sales on closed revenue, so the two are optimising different variables and the friction is rational. The best shared candidate in a long-cycle channel business is qualified pipeline created, in dollars, in territories where both functions operate. Marketing cannot inflate it without sales accepting the opportunities. Sales cannot claim it without marketing generating the input.

Then run a monthly meeting with one agenda item: a sample of ten rejected leads and ten won deals, read out loud, with both leaders present. Not a dashboard review, a file review. Disagreements that survive a dashboard rarely survive reading the actual records, because the arguments are nearly always about specific cases each side has generalised into a rule.

Closing: the discipline underneath all of this

Every KPI here lies the same way. It replaces a complicated reality with a single number, the number becomes the target, and people optimise the number instead of the reality. That is what metrics do. The defence is three habits.

Know who controls the definition. If the team measured on a number also controls what it counts, that number will improve and the improvement will be meaningless. That is the whole story of the MQL, and it applies equally to spec-in rate, co-op utilization, and trade show leads.

Match the measurement window to the buying cycle. In a nine to eighteen month capital purchase, most quarterly marketing metrics compare the wrong period's spend against the wrong period's results. Cutting investment improves nearly every efficiency metric for two to three quarters before the damage shows. That lag has ended more marketing budgets than poor performance ever has.

Prefer a rough measurement of the right thing to a precise measurement of the wrong one. A messy self-reported question on a quote form tells you more about a nine-month, dealer-mediated, twelve-person purchase than a perfectly instrumented last-click report ever will. Precision is not accuracy, and in long-cycle industrial marketing the most precise numbers on your dashboard are usually measuring the smallest and least important part of the journey.