Sales KPIs for Physical Goods and Channel Revenue

Fourteen sales metrics with named-source benchmarks, and how each one lies to you in a long-cycle channel business

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

Why almost every KPI article you have read does not apply to you

Search for sales KPIs and you get the same list written a hundred times: MRR, CAC payback, trial-to-paid, demo-to-close. That list was built by and for B2B SaaS. It assumes near-zero marginal cost, monthly billing, and a buyer who can say yes in a browser tab.

If you sell physical goods through dealers and distributors, none of that holds. Revenue arrives in lots. The decision to buy is governed by a harvest, a construction season, or a capital budget approved once a year. Your product occupies a yard, ties up floorplan financing, and depreciates while it waits. Freight can eat the margin your salesperson just negotiated. And your "customer" is often a dealer whose interests only partly overlap with yours.

Where a KPI behaves differently in the physical economy, I say so.

I took Ag Growth International's Brazil operation from roughly US$35 million to more than US$150 million in three years, as National Sales Manager, then General Manager Brazil, then Director North America. Grain and feed equipment, through dealers and distributors, OEM accounts and feedlots. Every number below is one I had to run, defend, or stop trusting.

1. Pipeline coverage

What it is

Pipeline coverage is the ratio of open opportunity value in a period to the revenue target for that same period.

The formula

Pipeline coverage = total open pipeline value for the period / quota for the period

Benchmark

No reliable public benchmark exists. The "3x rule" is folklore: no named research organization has published a study establishing it, and the arithmetic is circular, since 3x is the reciprocal of a 33 percent win rate. What is measurable is the gap between coverage and results. RAIN Group's 2026 B2B Sales Challenges Report found 47 percent of respondents reported an increase in qualified pipeline, while only 37 percent reported an increase in quota attainment (RAIN Group Center for Sales Research, 2026).

How this KPI lies to you

Coverage is the easiest number in the business to inflate, because the seller controls the numerator. Adding a stale opportunity costs nothing and instantly improves the ratio.

Seasonality compounds it. If revenue lands in two harvest windows, a quarterly ratio compares a full pipeline against a quota the season will not let you close. The same machine also sits in the pipeline twice, as a dealer stocking order and as an end-user order. And one speculative project can carry the whole ratio, because coverage measures presence, not probability.

What to do with it

Report three coverage numbers: pipeline created in the last 90 days, pipeline that moved a stage in the last 30, and everything else. The third bucket is usually the largest, and it is a graveyard. Set your target from your own trailing win rate.

2. Stage-by-stage conversion

What it is

Stage-by-stage conversion is the percentage of opportunities that advance from each stage to the next, and it is the only metric that shows where deals actually die.

The formula

Stage conversion = opportunities that reached stage N+1 / opportunities that entered stage N, on a fixed cohort created in the same window.

Benchmark

No reliable public benchmark exists, and none can: a firm with four stages and one with nine produce arithmetically incomparable rates. Only direction compares. Ebsta and Pavilion's 2025 GTM Benchmarks Report, built on 655,000 B2B opportunities and US$48 billion in pipeline, found top performers carry win rates 43 percent higher and sales cycles 42 percent shorter than bottom performers (Ebsta and Pavilion, 2025).

How this KPI lies to you

Stages are usually defined by what the seller has done, not what the buyer has done. "Proposal sent" is a seller action; "budget confirmed by the person who signs" is a buyer action. When stages track seller activity, conversion measures diligence, not progress, and a rep can walk a dead deal through five stages. In channel sales the deal also lives inside the dealer's process, which your CRM cannot see: your funnel shows "quoted" while the dealer has not yet called the farmer.

What to do with it

Rewrite every stage as a buyer-verifiable event, and measure conversion by creation month rather than close month. In a channel, add one stage only the dealer can confirm, such as "end user identified and named."

3. Sales cycle length

What it is

Sales cycle length is the elapsed time from creation of a qualified opportunity to won or lost, and it sets the lead time between a change today and the revenue it produces.

The formula

Sales cycle length = close date minus qualified-opportunity creation date, reported as a median and by deal-size band, never as a single mean.

Benchmark

No reliable public benchmark exists for capital equipment or channel-based physical goods; republished "average B2B sales cycle" figures trace back to SaaS datasets. Direction is documented: RAIN Group's 2026 B2B Sales Challenges Report found cycle length increased for 48 percent of respondents (RAIN Group, 2026). The United States Department of Agriculture Economic Research Service forecast 2026 net farm income at US$153.4 billion, 0.7 percent below 2025 (USDA ERS, 2026). When farm margins compress, equipment decisions get postponed and cycles stretch for reasons no process fix will reach.

How this KPI lies to you

The mean is dominated by deals that closed, and closed deals are disproportionately the fast ones. Losses and no-decisions sit open longer, and many are never formally closed, so they never enter the calculation.

The same product sold in February and in August has different cycle lengths, because the February buyer has a planting deadline and the August buyer does not. The annual average describes no actual month.

What to do with it

Report the median by season and deal-size band, including lost and no-decision opportunities. Then use it for one thing: the cutoff date after which new pipeline cannot convert this fiscal year.

4. Average deal size

What it is

Average deal size, or average ticket, is mean revenue per closed-won order, and it is the metric most likely to be a statistical artifact in a physical-goods business.

The formula

Average deal size = total closed-won revenue in the period / number of closed-won orders in the period

Benchmark

No cross-industry benchmark exists. The useful reference is a contrast: The Bridge Group's 2026 research across 158 B2B companies puts the median annual quota for a software account executive at US$960,000 against median on-target earnings of US$200,000, a ratio of 4.6 to 1 (The Bridge Group, 2026). Those economics belong to near-zero cost of goods; at equipment margins in the twenties or thirties that ratio would bankrupt you.

How this KPI lies to you

Deal size distributions in capital goods are severely right-skewed. One turnkey grain facility outweighs two hundred small orders, so the mean sits above the great majority of your transactions. Median and mean can differ by a factor of five in the same book.

It is also gameable without lying. Bundle freight, installation, parts and warranty into one line and average ticket rises with no change in customer behavior, often at lower margin. In a devaluing market it rises by itself, and management congratulates a sales team for the exchange rate.

What to do with it

Report the median alongside the mean and alongside unit volume. If ticket rises while units fall, you are losing the small accounts that feed parts and service revenue for a decade. Track it in constant currency, net of freight.

5. Win rate

What it is

Win rate is the share of qualified opportunities that end in a signed order, and it is the most manipulable number in sales because the seller controls the denominator.

The formula

Win rate = opportunities won / (won + lost + no decision)

The third term gets quietly dropped, and dropping it turns a 20 percent win rate into 35 percent without a single extra order.

Benchmark

Ebsta and Pavilion's 2025 GTM Benchmarks Report, covering 655,000 opportunities and more than 2,000 revenue leaders, found win rates still falling year over year, though less steeply: down 18 percent in 2024 and 10 percent in 2025. Top performers close 11 times faster than bottom performers, a gap that widened from 8.9 times in 2024 (Ebsta and Pavilion, 2025). Absolute win rate levels for channel-based capital equipment have no reliable public benchmark.

How this KPI lies to you

Win rate rewards not competing. A seller who logs only the opportunities they expect to win posts the best win rate on the team and the worst pipeline contribution. That is the predictable response to paying on it.

In channel sales you often do not know you competed. The dealer quoted three brands and presented one. Your CRM records a clean win and no loss, so your win rate is a record of the deals your dealer chose to show you.

What to do with it

Split win rate into three permanent numbers: new accounts, competitive replacements, repeat business. Force every opportunity to a terminal status, with "no decision" explicit. Never pay on win rate; pay on won revenue and margin.

6. Quota attainment

What it is

Quota attainment is the percentage of assigned target a seller delivers, and its distribution says more about how you set quotas than about how your people sell.

The formula

Quota attainment = actual revenue (or margin) delivered / assigned quota for the period. The organizational metric that matters is percentage of sellers at or above 100 percent.

Benchmark

The Bridge Group's 2026 research across 158 B2B companies found 48 percent of account executives reached quota (The Bridge Group, 2026). Ebsta and Pavilion found 78 percent of sellers missed quota in 2025, up from 69 percent in 2024 (Ebsta and Pavilion, 2025). Both are technology-weighted: evidence that quota-setting is broadly broken, not a target for an equipment business.

How this KPI lies to you

Average attainment is close to useless, because the distribution is what matters. A team averaging 100 percent can be one seller at 250 and six between 55 and 70: a single point of failure with an org chart attached.

Attainment is also quota quality disguised as performance. If most of the team misses, the likeliest cause is quotas set from a budget rather than territory potential. Physical businesses worsen this by splitting an annual target into equal monthly quotas in a market with two selling seasons, guaranteeing months where nobody can hit the number.

What to do with it

Publish the distribution: percentage above 100, between 80 and 100, below 60. If fewer than half clear 100, fix the quota model before you touch the team. Seasonalize quotas, and set them from measured potential: installed base, planted area, herd count.

7. Seller ramp time

What it is

Seller ramp time is the number of months from a new seller's start date until they produce consistently at the level expected of a tenured seller.

The formula

Ramp time = months from start date to the first of three consecutive months at or above 100 percent of a fully-ramped monthly quota

Benchmark

The Bridge Group's 2026 research across 158 B2B companies reports ramp time at 6.2 months, the highest in the study's history, even though companies now hire sellers averaging 3.7 years of experience (The Bridge Group, 2026). For capital equipment sold through a dealer network, no reliable public benchmark exists. My working assumption is one full sales cycle plus one full season, because a seller cannot be judged until they have carried deals through a complete demand cycle: a year or more, roughly double the software figure.

How this KPI lies to you

Ramp is measured against a ramped quota that was itself set optimistically, so a shortening ramp can simply mean quotas got easier. Worse, new sellers inherit accounts with orders already in motion, and those close under the new name, so the curve measures territory inheritance, not capability.

A seller starting three months before peak season ramps two or three times faster than an identical seller starting the week after it ends. Compare them and you are ranking start dates.

What to do with it

Measure ramp only on opportunities created after the start date, and report the start month alongside it. Use it as a planning constraint: if ramp is nine months, a seller hired in the second half contributes nothing this year.

8. Forecast accuracy

What it is

Forecast accuracy is the difference between what you committed to deliver at the start of a period and what you delivered, expressed as a percentage of the commitment.

The formula

Forecast accuracy = 1 minus (absolute value of (actual minus forecast) / forecast)

The absolute value matters. Without it, a quarter 20 percent over and a quarter 20 percent under cancel out and the annual average looks excellent.

Benchmark

I could not verify a forecast accuracy benchmark at a primary source. The most circulated figure, that fewer than one in four sales organizations forecast within 10 percent of actuals, is routinely attributed to Gartner, but it does not appear in Gartner's own public material, so I will not present it as established. The nearest verifiable proxy is quota attainment: Ebsta and Pavilion found 78 percent of sellers missed quota in 2025. If most sellers miss, the forecast built from their inputs was wrong.

How this KPI lies to you

Signed error cancels. A team alternating over and under looks accurate annually while being unmanageable month to month, and freight and working capital are planned monthly.

The larger deception is that accuracy improves when you sandbag. An organization that learns accuracy is the compensated metric will commit low, deliver above, and post excellent accuracy while understating capacity. That is expensive in a way it never is in software: you build to the forecast, so understating it means missed builds, expedited freight, and lost orders you never see.

What to do with it

Measure three things and never blend them: accuracy of orders booked, of revenue shipped, and of margin. Track bias separately as signed average error, because it tells you whether the team sandbags or dreams.

9. Revenue per seller

What it is

Revenue per seller is total revenue divided by quota-carrying headcount, and it is the crudest productivity measure and the one most often used to justify the wrong decision.

The formula

Revenue per seller = total revenue in the period / average number of quota-carrying sellers. In physical goods, run it again as gross margin per seller, which is the version that pays salaries.

Benchmark

The Bridge Group's 2026 median software account executive quota of US$960,000 (The Bridge Group, 2026) is the wrong reference class for equipment. For distribution the named source is the North American Equipment Dealers Association Cost of Doing Business Study. Its 2025 edition, reporting 2024 results as summarized by Farm Equipment, shows parts gross margin of 30.06 percent for dealers in the US$200 million to US$600 million band and 26.98 percent above US$600 million, with service gross margin of 67.38 percent and 52.18 percent (NAEDA, 2025).

How this KPI lies to you

Revenue per seller rises fastest when you cut sellers, which is why it is dangerous. Remove your two weakest territories and the ratio improves immediately, while the revenue they would have produced in three years disappears.

It also fails to separate selling from order-taking, since much of a mature network's volume is repeat business that would arrive without a salesperson. And measured on revenue rather than margin, it rewards discounting, because the fastest way to raise the number is to cut the price.

What to do with it

Report gross margin per seller net of freight, split into accounts that existed at the start of the year and accounts that did not. The second number is what your sales team actually produced.

10. Discount rate and margin erosion

What it is

Discount rate is the gap between list price and realized price, and margin erosion is what that gap does to gross profit once freight, rebates, program money and terms are counted.

The formula

Discount rate = 1 minus (net realized price / list price)

Pocket margin = net realized price minus cost of goods minus freight minus rebates minus program costs minus the financing cost of terms granted

Benchmark

Simon-Kucher's Global Pricing Study 2025, based on more than 2,200 business leaders across 28 countries and 39 industries, found companies realize less than half of the price increases they announce, with average price realization down 5 percentage points over two years (Simon-Kucher, 2025). The NAEDA 2025 Cost of Doing Business Study (2024 results) reported return on equity falling from 30.93 percent to 21.59 percent for dealers in the US$200 million to US$600 million band, and from 32.82 percent to 14.04 percent above US$600 million (NAEDA, 2025).

How this KPI lies to you

Reported discount is measured at the invoice, and the invoice is the smallest part of the giveaway. Extended terms, absorbed freight, extra warranty, floorplan support, co-op marketing and volume rebates never appear in the discount field. I have seen a 4 percent invoice discount accompanied by more than twice that value in off-invoice support.

Averaging is the second failure: a 6 percent average can be most orders at zero and a handful at 25 percent, and the handful are your largest accounts, so your worst margins sit on your biggest volume. The third lie is timing: trade-ins and buy-backs convert a concession today into an inventory loss two years from now, attributed to nobody.

What to do with it

Build a pocket-margin waterfall from list down to actual contribution, with every off-invoice item on it, and report discount at pocket level only. Then rank your top twenty accounts by pocket margin. In most physical-goods businesses at least one of the top five by revenue sits in the bottom five by contribution.

11. Activity versus outcome

What it is

Activity metrics count what sellers do (calls, visits, quotes, demonstrations), outcome metrics count what results, and the ratio between them is useful only when it is stable.

The formula

Activity efficiency = outcomes produced / activities performed, for example orders per dealer visit or orders per quote issued.

Benchmark

No reliable public benchmark exists for the correct volume of sales activity, and any number offered is guesswork dressed as research, because the right level depends on territory density, cycle length and product complexity. What is documented is how little of a seller's time reaches the customer: Salesforce research published in 2023 found representatives spend less than 30 percent of their time selling (Salesforce, 2023).

How this KPI lies to you

Any activity metric that is managed becomes a target, and any target becomes easy to satisfy without doing the work. Require twelve dealer visits a month and you get twelve visits, several of them coffee with the friendly dealers nearest the office. The metric looks healthy right up to the moment the territory stops growing.

Field selling also carries windshield time. A seller covering three states may spend two days a week driving, so counting visits without counting distance rewards whoever has the densest territory, which is usually your best-covered market and smallest growth opportunity.

What to do with it

Never set activity targets in isolation. Track the ratio and investigate it when it moves, not the raw count. Weight activity by account potential, so a dealer with US$5 million of potential is not scored like one with US$300,000.

12. Sell-in versus sell-through

What it is

Sell-in is what you invoice to your dealers, sell-through is what those dealers invoice to the end user, and the difference is inventory sitting in your channel that you already booked as revenue.

The formula

Sell-in = units or value shipped to channel partners in the period

Sell-through = units or value sold by channel partners to end users in the period

Channel inventory change = sell-in minus sell-through

Benchmark

No reliable cross-industry benchmark exists for sell-through rate. The figures online (commonly "80 percent or above") come from retail merchandising, where product turns in weeks. Two named public sources measure channel health directly. The United States Census Bureau's Manufacturers' Shipments, Inventories and Orders survey shows the machinery inventories-to-shipments ratio, not seasonally adjusted, moving between 2.25 and 2.91 in the eight months to June 2026, against total manufacturing at 1.48 seasonally adjusted (Census Bureau M3 survey, 2026, via FRED). The Association of Equipment Manufacturers publishes monthly United States retail unit sales: March 2026 tractors down 9.1 percent, combines down 25.3 percent against March 2025 (AEM, 2026).

How this KPI lies to you

Sell-in is the number in your P&L, and it can grow for two full quarters while end-user demand falls, because you are filling dealer yards rather than farms. This is the most expensive lie in channel sales, and it self-corrects in the worst way: a quarter with almost no orders, because the channel is full and no dealer will stock another unit. Incentives make it deliberate, since volume rebates pay dealers to buy, not to sell.

Sell-through lies a second way: it is self-reported by partners with an interest in the answer. A dealer chasing program money reports optimistically; one negotiating floorplan relief reports pessimistically.

What to do with it

Make sell-through reporting a condition of the dealer agreement, at serial number level, and reconcile it against warranty registrations, which the end user files independently. Then shift part of dealer incentive from purchase volume to registered end-user sales.

13. Territory coverage

What it is

Territory coverage is the share of addressable opportunity inside a geography that your channel and sales team actually reach, and it reveals growth you are not pursuing rather than growth you are losing.

The formula

Territory coverage = accounts actively worked in the period / total qualified accounts in the territory

Territory share = your sales in the territory / total addressable demand, using an external denominator such as registered units, planted area or herd numbers.

Benchmark

No reliable public benchmark exists. A repeated claim that roughly 55 percent of sales territories are misaligned is attributed to Andris Zoltners and colleagues, but I could not verify it at a primary source and do not present it as established. The underlying work is real: Zoltners, Sinha and Lorimer argue in Harvard Business Review (2015) that territory design remains one of the most undervalued levers in sales productivity.

How this KPI lies to you

Coverage is normally computed against the accounts already in your CRM, so it measures how well you cover the accounts you already know about. A territory can show 92 percent coverage while a third of the real buyers there were never entered. The metric reports confidently on a denominator you invented.

In dealer networks, coverage is assumed rather than measured. A dealer holds a contract for a region, the map is colored in, headquarters records it as covered. The dealer may work 40 percent of it, and that white space is invisible precisely because it is contractually assigned.

What to do with it

Build the denominator outside your CRM, using registration data, association statistics or agricultural census figures, and compare it against your active account list. The gap is your real white space. Report coverage of potential, not coverage of accounts.

14. Channel inventory turns and aging

What it is

Channel inventory turns measure how many times a dealer sells through stock in a year, and aging measures how long each unit has been sitting.

The formula

Inventory turns = cost of goods sold in the period / average inventory at cost

Aging = share of channel units on hand more than 6, 12 and 18 months

Benchmark

The North American Equipment Dealers Association Cost of Doing Business Study, as summarized by Farm Equipment, which also notes an operating target above 3 turns and a North American average near 2.8.

Metric	US$200M to US$600M dealers	Above US$600M dealers
Inventory turns, 2025 edition (2024 results)	2.52	2.82
Inventory turns, 2024 edition (2023 results)	2.82	3.16
Parts aftermarket absorption, 2025 edition	38.45%	26.68%
Parts aftermarket absorption, 2024 edition	40.02%	35.95%

How this KPI lies to you

Turns are a ratio of a flow to an average, and the average smooths exactly the problem you are looking for. A dealer can post acceptable annual turns while holding four units that have not moved in two years, because parts volume carries the ratio.

Turns also improve when a dealer stops buying: inventory falls, the ratio rises, and the dashboard reports improvement in the exact quarter the channel stopped stocking your product. And because turns are computed at cost, a dealer who marks down old units shows better turns and worse profit at once, so the metric applauds a loss.

What to do with it

Track turns and aging together, and set stocking programs against aging rather than volume. Review the aged tail by serial number in every dealer business review, and fund clearance before units cross twelve months. A dealer carrying an aged tail will not take new product whatever you offer.

The KPIs you should stop tracking

- Number of leads. A volume metric with no connection to revenue in a market with a finite, known set of buyers. If you sell grain handling equipment you can name most qualified buyers in a region. Replace it with named-account penetration.
- Blended win rate. It cannot be acted on, because it mixes repeat orders that were never at risk with competitive fights you lost. Keep the segmented versions and delete the blended one.
- Average deal size as a headline metric. In a right-skewed, bundled, currency-exposed business it moves for accounting reasons more often than commercial ones. Keep it in analysis, off the board slide.
- Total pipeline value. The number most likely to be inflated and least likely to be questioned, because leaving a dead opportunity open costs nothing. Replace it with the two movement measures above.
- Activity counts as standalone targets. Count activities to compute ratios, never to set targets. A target on a count guarantees the count and nothing else.

What to measure in your first 30 days

Walking into an organization with no reliable measurement, do not attempt the full set. Six numbers, in this order.

1. Revenue and gross margin by account, last 24 months. By account, not by region or product line, and 24 months so you see one seasonal cycle twice. This usually reveals concentration risk and at least one large account served at a loss.
2. Sell-in versus sell-through by dealer, last 12 months. If sell-through data does not exist, that is the most important thing you will learn in month one.
3. Pipeline aging. Every open opportunity by date of last stage movement. Expect a large share to have sat still over 90 days. Reconciling it is a two-week job that changes the forecast immediately.
4. Quota attainment distribution, not average. Percentage of sellers above 100, between 80 and 100, and below 60. One page tells you whether you have a talent problem or a quota problem.
5. Pocket margin on the top 20 accounts. List price down to actual contribution. Budget a week, because the data is rarely in one system.
6. Forecast versus actual for the last four periods, by absolute error and by bias. Absolute error tells you how good the forecast is. Signed bias tells you whether the organization sandbags or dreams.

Notice what is absent: activity metrics, lead counts, and anything requiring a new system. All six come from data that already exists.

A closing note on measurement

The pattern running through all fourteen KPIs is that every metric can be satisfied without the business improving. Coverage rises when you leave dead deals open. Win rate rises when you compete less. Turns improve when the channel stops buying. Accuracy improves when you commit low.

That is a property of measurement, not cynicism about salespeople. The defense is not more metrics. It is pairing every metric with the one that moves the opposite way when the gaming starts: coverage with movement, revenue with margin, sell-in with sell-through, turns with aging, accuracy with bias. In a seasonal business, give it two full cycles before you trust a conclusion.

Andre Magrini is a fractional Chief Revenue Officer based in Chicago. He led Ag Growth International's Brazil operation from roughly US$35 million to more than US$150 million in three years as National Sales Manager, then General Manager Brazil, then Director North America covering the United States and Canada, selling grain and feed equipment through dealers and distributors, OEM accounts and feedlots.