

Customer Service KPIs for Equipment and Industrial Firms

Nineteen service metrics, how each one lies to you, and what to measure when support runs through the dealer

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

By Andre Magrini, fractional Chief Revenue Officer, Chicago. Andre led the Brazil operation of Ag Growth International from roughly US\$35 million to more than US\$150 million in three years, selling through dealers, OEM accounts and feedlots.

Why most support KPI guides do not apply to you

Almost everything published about customer service metrics was written for a software help desk. The unit of work is a ticket, the customer is a user with a login, and the win is deflection: pushing him to an article so no human is involved. That model breaks when your product weighs four tons and sits in a field. A support event in an equipment business is a grain leg that stopped at 2am with trucks lined up, or a technician three hours out who brought the wrong seal kit.

Support is not a cost center, it is the sales channel. In software, bad support costs you a renewal that is already priced. In equipment, it costs you the next machine and the dealer's willingness to keep your line on his lot.

Deflection is usually the wrong goal. A farmer who fixes his own problem at 2am with a video is a win. A farmer who gives up trying to reach you and calls the aftermarket jobber is a loss. Both appear in your data as "no contact." In equipment the alternative to contacting you is someone else, and that someone else is now in your account.

Each KPI below gets five things: what it is, the formula, a benchmark with a named source (or an honest statement that none exists), how it lies to you, and what to do with it.

Part 1: satisfaction and loyalty

1. Customer Satisfaction Score (CSAT)

What it is. The percentage of customers who rate a specific service interaction as satisfactory.

Formula. $CSAT = (\text{satisfied responses, typically 4 and 5 on a 5 point scale}) / (\text{total responses}) \times 100$.

Benchmark. The American Customer Satisfaction Index reported a national score of 76.9 out of 100 for the fourth quarter of 2025, and noted the index has not materially increased since 2017. Closer to durable goods, the ACSI Household Appliance and Electronics Study 2025 put household appliances at 81. MetricNet's contact center quartiles place top quartile CSAT at 94.6% or higher and bottom quartile below 72.6%.

How it lies to you. It surveys only closed interactions, so the customer who called twice, got voicemail and called your competitor's dealer is never in the denominator. And it measures the interaction, not the outcome: a customer can rate a call 5 out of 5 while the machine is still down. Politeness scores well. Uptime is what he bought.

What to do with it. Survey on outcome, after the machine is running, report response rate beside the score, and call a rotating sample of accounts that filed no support event last quarter to ask who fixed what broke.

2. Net Promoter Score (NPS), and the honest critique

What it is. A loyalty index built on one question: how likely is the customer to recommend you, scored 0 to 10.

Formula. $NPS = (\text{percent scoring 9 or 10}) - (\text{percent scoring 0 to 6})$.

Benchmark. Qualtrics XM Institute's 2024 XMI Customer Ratings for NPS covered 22 industries, with grocers highest at 34.3 and car rental lowest at 15.8. No comparable published benchmark exists for equipment manufacturers.

How it lies to you. Its central claim did not survive testing. Keiningham, Cooil, Andreassen and Aksoy, in the Journal of Marketing (July 2007), used data from 21 firms and more than 15,500 interviews and could not replicate the claimed superiority of Net Promoter over other satisfaction measures, including in the industries cited as its exemplars. In a dealer channel the question is also ambiguous: a farmer answers about his dealer, his salesman and his last repair bill, in proportions you cannot recover. And with 40 respondents out of 300 accounts, a handful of detractors swings the score by double digits.

What to do with it. Keep it if your board reads it, but demote it to directional, because the value sits in the free-text follow-up. If you want one loyalty number that predicts behavior, use repurchase rate (KPI 18): it is observed, not stated.

3. Customer Effort Score (CES)

What it is. A measure of how much work the customer had to do to get his problem solved.

Formula. Agreement with a statement such as "the company made it easy for me to handle my issue," on a 1 to 7 scale, reported as the mean or as top-two-box percentage.

Benchmark. CES was developed in 2010 by Matthew Dixon, Karen Freeman and Nick Toman at the Corporate Executive Board (now Gartner). Their Harvard Business Review article "Stop Trying to Delight Your Customers" (July/August 2010), drawing on more than 75,000 customer interactions, found that 96% of customers with a high-effort service interaction became more disloyal, against 9% with a low-effort one. No cross-industry CES benchmark is published.

How it lies to you. It measures the effort the customer noticed. A customer whose rep quietly found the part in another branch reports low effort while your system was failing expensively behind the curtain. Effort is also unevenly distributed: the owner answers your survey, the mechanic who spent five hours on the phone does not.

What to do with it. This is the best of the three attitudinal metrics for equipment, because effort maps onto what the customer is losing, which is time. Ask it after high-severity events only, then translate it into hours, because "our customers average X hours from breakdown to running" is a CES statement you can put in a bid.

Part 2: resolution

4. First Contact Resolution (FCR)

What it is. The share of issues fully resolved in a single contact, with no callback, transfer or second visit.

Formula. $FCR = (\text{issues resolved on first contact}) / (\text{total issues}) \times 100$.

Benchmark. SQM Group's 2025 research puts the aggregated all-industry average at 70%, with an industry range of roughly 50% to 90% by call complexity. MetricNet places top-quartile FCR at 86.3% or higher.

How it lies to you. This is the most gamed metric in support, because "resolved" is defined by the person being measured: a case closed today and reopened Thursday records as two resolved contacts rather than one failure. It also penalizes taking hard problems, since a team measured on FCR routes anything complicated into a specialist queue, producing a beautiful front-line number and a backlog of exactly the issues that lose customers.

What to do with it. Redefine it in writing first: resolved only if no contact about the same asset and the same failure mode arrives within 14 days. Then pay on FCR and reopen rate together, never on FCR alone, so closing without solving costs the same person twice.

5. First Response Time

What it is. Elapsed time between the customer raising an issue and a human acknowledging it substantively.

Formula. Median and 90th percentile of (first human response minus customer contact).

Benchmark. Zendesk's CX Trends framing treats 12 hours or less as good, 4 hours or less as better and 1 hour or less as best for email. MetricNet puts top-quartile average speed of answer at 29 seconds or less for voice. Treat cross-industry averages here with suspicion, because targets vary so much by channel and contract tier.

How it lies to you. Averages hide the cases that cost you the account: a 40 minute mean can sit on an 11 hour 90th percentile, and the 11 hour cases cause the churn. The clock also starts at the wrong zero: yours when the ticket is created, his when the machine stopped.

What to do with it. Split by severity and by season, because four hours in February and four hours during harvest are not the same product. Set a hard commitment for machine-down events only, staff to hit it at the 90th percentile, and publish it: a fast response you do not promise earns you nothing.

6. Time to Resolution

What it is. Total elapsed time from the customer raising the issue to genuine closure.

Formula. Median and 90th percentile of (resolution timestamp minus contact timestamp).

Benchmark. No credible cross-industry benchmark exists for equipment support, and this guide will not manufacture one: the figures circulating online come from software help desk datasets where resolution does not involve shipping a physical object.

How it lies to you. Waiting time gets excluded. Teams pause the clock for the customer, for a part, for a third party, each pause defensible alone, and together they produce a number with no relationship to being down. It also rewards closing over fixing: the fastest path to a good number is to close the case, which returns next week with a fresh clock.

What to do with it. Run two clocks: the internal one you can pause shows where your process is slow, and the customer one (breakdown to running) shows what you are selling. The gap between them is your parts and logistics problem, not your support problem.

7. Reopen Rate

What it is. The share of closed cases reopened, or that generate a new case for the same problem on the same asset.

Formula. Reopen Rate = (cases reopened or duplicated within the window) divided by (cases closed in the period), times 100.

Benchmark. No authoritative cross-industry benchmark exists. The figures circulating (commonly "under 5% is healthy") come from vendor blogs and single-firm samples in managed IT services, not from a research body.

How it lies to you. It is trivially suppressed: if your process opens a new case instead of reopening the old one, the rate is structurally near zero and uninformative. It also undercounts the worst outcome, because the customer who gives up generates no reopen at all.

What to do with it. Link cases by serial number and failure code rather than by ticket ID, so a "new" case on the same machine for the same fault is flagged automatically. Repeat failures on one serial number are a product problem in a support costume.

8. SLA Adherence

What it is. The percentage of cases handled within the response and resolution commitments you contractually promised.

Formula. SLA Adherence = (cases meeting the committed threshold) divided by (total cases subject to that SLA), times 100.

Benchmark. Aberdeen Group's March 2016 study "First-Time Fix: A Metric That Drives Success" (Aly Pinder Jr., n=182) reported SLA compliance of 85% among organizations with a first-time fix rate above 70%, against 74% among those at or below it. On the target itself, the ubiquitous 80/20 standard (80% of calls answered within 20 seconds) has no published research basis; its origin is disputed, attributed to early AT&T practice or to Rockwell's 1970s call center hardware. It survived because it was easy to write into a contract.

How it lies to you. You wrote the test, and teams under pressure renegotiate the target rather than improve. It measures the timer rather than the outcome, so meeting a four hour on-site commitment and failing to fix the machine scores as compliant. Aberdeen's own recommendation was that service teams be measured on the customer's outcome instead.

What to do with it. Keep it as contractual hygiene, not as performance, and audit the exclusion list annually. If the share of volume falling outside SLA is growing, your adherence improvement is fiction.

Part 3: load and cost

9. Contact Volume per Customer

What it is. How many support events an account generates, normalized by that account's installed base.

Formula. Contacts per Unit = (support events from the account) divided by (units in service at that account), per period.

Benchmark. No credible public benchmark exists; any figure quoted as one is a software SaaS average.

How it lies to you. High volume gets read as a problem account. It is frequently the opposite: it is the account that trusts you enough to call. The dangerous accounts are the quiet ones with a large installed base, because their machines are failing at the population rate and someone else is fixing them.

What to do with it. Rank accounts by contacts per unit and look at both tails: the high tail shows where the product or the training is failing, and the low tail is a prospecting list for your aftermarket team and your best early warning of a lost account.

10. Contact Volume per Product

What it is. Support events attributed to a specific model, configuration, production batch or component.

Formula. $\text{Contacts per Unit in Field} = (\text{support events for that model}) / (\text{units of that model in service})$. Normalize by operating hours where telematics allow.

Benchmark. No cross-industry benchmark exists. What matters is the difference between your own models and batches.

How it lies to you. Raw counts follow sales volume, so your best-selling model always looks like your worst-quality model. Without normalizing by units in service, and ideally by hours, this metric points engineering at exactly the wrong product.

What to do with it. Route it into design review tagged by serial number, failure code and production batch, so an anomalous batch becomes a containment decision while it is still cheap. In most equipment companies this is the highest-value support metric and the least connected to engineering.

11. Deflection Rate

What it is. The share of issues resolved through self-service without a human on your side.

Formula. $\text{Deflection Rate} = (\text{self-service sessions with no subsequent contact}) / (\text{total self-service sessions}) \times 100$.

Benchmark. Gartner's press release of 19 August 2024, based on a survey of 5,728 customers, found that only 14% of customer service issues are fully resolved in self-service, rising to just 36% for issues customers called very simple. Adoption is high (73% use self-service at some point). Resolution is not.

How it lies to you. This is the most dangerous metric in the guide for an equipment business, because the measurement is an absence. Abandonment and resolution look identical: the customer did not contact you. In software that means he gave up. In equipment it means he called the aftermarket supplier or your competitor's dealer, so your deflection number improves while your parts revenue leaves.

What to do with it. Stop targeting it, because a deflection target changes behavior in one direction only: it makes reaching a human harder. Measure self-service like a product (did he complete the task and confirm it) and keep the path to a human on machine-down events at one step.

12. Cost per Contact

What it is. Fully loaded support operating cost divided by support events handled.

Formula. $\text{Cost per Contact} = (\text{support operating expense: salaries and benefits for agents and indirect staff, technology, telecom, facilities, travel and training}) / (\text{contacts handled})$.

Benchmark. MetricNet, in Jeff Rumburg's "Metric of the Month: Cost per Ticket" published through HDI in 2011 using 2010 North American data, reported an average desktop support cost per ticket of US\$62 with a range from US\$27 to US\$490. That data is dated and comes from IT support, so use the dollar figure cautiously. What has not aged is the spread: unit costs varied by more than a factor of ten across comparable organizations.

How it lies to you. The denominator is the easiest thing in the company to move: split one issue into three cases and unit cost falls by two thirds with nothing changed. It is also blind to value, averaging a US\$400 call that saves a US\$600,000 relationship with a US\$40 call about a manual.

What to do with it. Never set a reduction target on it. Use cost per resolved issue instead, which removes the case-splitting incentive, and pair it with the full cost of the failure.

Part 4: the metrics that decide equipment businesses

These separate an equipment support organization from a help desk. Most published KPI guides do not mention any of them.

13. Customer Equipment Uptime and Downtime

What it is. The percentage of scheduled operating time during which the customer's machine was available to run.

Formula. $\text{Uptime} = (\text{scheduled operating hours} - \text{unplanned downtime hours}) / (\text{scheduled operating hours}) \times 100$.

Benchmark. Aberdeen Group (March 2016, n=182) reported serviceable asset uptime of 84% among service organizations with a first-time fix rate above 70%, against 71% among those at or below it. The Senseye and Siemens report The True Cost of Downtime 2024 estimated that unplanned downtime costs the world's 500 largest companies around US\$1.4 trillion a year, about 11% of turnover, up from 8% in 2019. For agriculture, the U.S. PIRG Education Fund and National Farmers Union report "Out to Pasture" (April 2023), surveying 53 farmers in 14 states, modeled average losses of US\$3,348 per farmer per year from repair-restriction downtime (small sample, advocacy publication, and this guide flags both).

How it lies to you. The denominator is negotiable. "Scheduled operating hours" excludes time the machine was not scheduled to run, and in a seasonal business that is the whole game: a combine down six days in January costs almost nothing, and six hours in October can cost the crop. Uptime is also jointly produced by your machine, the dealer's technician, the customer's maintenance discipline and your parts network, then reported as one number that says nothing about where to intervene.

What to do with it. Measure uptime inside the customer's critical window, not across the calendar. Define that window with him (planting, harvest, a production campaign, a feedlot's schedule) and report in-window uptime as the headline.

14. Downtime Hours per Unit, by Cause

What it is. Absolute downtime hours per machine per period, split by what the customer was waiting on.

Formula. $(\text{Total unplanned downtime hours across the fleet}) / (\text{units in service})$, broken into time to diagnosis, time to part, time to technician and time to authorization.

Benchmark. No public benchmark exists, because almost nobody publishes this decomposition. Your own baseline is the benchmark, and building it is the exercise.

How it lies to you. Only if you fail to decompose it. An aggregate downtime number invites everyone to blame someone else and produces no decision. Split, it becomes uncomfortable and useful: in most equipment companies the largest bucket turns out to be time to part, followed by time to authorization, and neither is a support department problem.

What to do with it. Present the four buckets quarterly and let whichever is largest own the improvement budget. This one chart moves more revenue than any survey score in this guide.

15. Time to Part on the Ground, and Parts Fill Rate

What it is. How long the customer waits for the part that will fix his machine, and how often you had it in stock when he asked.

Formula. Time to Part = median and 90th percentile of (part arrival at the machine minus part identification). Fill Rate = (order lines shipped complete from stock on first request) divided by (total order lines), times 100.

Benchmark. No independently audited cross-industry fill rate benchmark exists for agricultural or heavy equipment. Figures such as "90% off-the-shelf fill rate" circulate widely and are attributed to automotive dealer associations, but this guide could not verify a primary source and will not present them as a benchmark. What can be sourced is the consequence: Aberdeen (2016) found the number one reason customers complain about field service is that the technician could not resolve the issue, cited by 61% of respondents, against 51% for appointment waiting time and 41% for late arrival.

How it lies to you. Time to part is usually measured dock to dock, which looks excellent while the customer waits, because his clock starts when he calls and ends when the machine runs. Fill rate is worse: it is measured against what was ordered, not what was needed, so a customer who already knows you do not stock a part orders it elsewhere and never enters your denominator. Chronic stockouts eventually produce excellent fill rates. The aggregate is also carried by filters and belts, while the slow-moving part that stops a machine is a rounding error in the average.

What to do with it. Measure door to machine on the customer's clock and publish the 90th percentile. Build a short list of machine-down parts per model and measure fill rate on that list alone: that number predicts your next machine sale.

16. Warranty Cost as a Percentage of Revenue

What it is. What you spend honoring warranty obligations, against the product revenue that generated them.

Formula. Warranty Claims Rate = (warranty claims paid) divided by (product revenue), times 100. Track the accrual rate (what you set aside) separately from the claims rate (what you paid).

Benchmark. Warranty Week's 23rd Annual Product Warranty Report (16 April 2026) reported that in 2025 U.S. based manufacturers ran an average warranty claims rate of 1.30% of product sales and an average accrual rate of 1.43%. Across the 23 year series from 2003, the average claims rate was 1.42% with a standard deviation of 0.18%. The vehicle sector, which includes trucks and buses, accounted for 65% of total accruals.

How it lies to you. The ratio moves for reasons unrelated to quality: a strong sales year inflates the denominator while last year's production is still failing, because warranty cost lags revenue by the term. It is

also the easiest number in the company to manage down without improving anything. Tighten adjudication, push decisions to the dealer, narrow coverage, and the ratio falls. You transferred cost onto the dealer, and you will pay for it in dealer renewal instead.

What to do with it. Track claims cost by production cohort against machine hours, not by calendar period against current revenue. Then measure the number everyone avoids: warranty claim cycle time from dealer submission to payment. Slow payment is the best predictor of a dealer moving your line to the back of the lot.

17. First-Time Fix Rate

What it is. The share of field service visits that fully resolve the issue on the first trip.

Formula. $FTFR = (\text{work orders resolved on first visit}) / (\text{total work orders})$, times 100. Any repeat visit or secondary truck roll for the same issue disqualifies the original.

Benchmark. Aberdeen Group's March 2016 study used a 70% threshold to split its sample. Organizations above 70% FTFR reported customer retention of 86%, SLA compliance 85%, asset uptime 84%, satisfaction 85% and service revenue improving 4% year over year. Those at or below 70% reported retention 76%, SLA compliance 74%, uptime 71%, satisfaction 66% and no change in service revenue. The study is dated, but it remains the most cleanly documented public link between first-time fix and commercial outcomes.

How it lies to you. "Fixed" is recorded by the technician who wants to leave, so unless a second visit for the same fault on the same serial number is linked to the first, FTFR measures technician optimism. It also punishes the technician for the parts department's failure, since the most common cause of a failed first visit is the wrong part on the truck.

What to do with it. Report FTFR and time to resolution side by side permanently, because each constrains the gaming of the other. Attribute every failed first visit to a cause code (part not available, wrong diagnosis, skill gap, site access) and manage the causes, not the rate.

The KPI that ties support to revenue

Everything above measures the support function. None of it measures why the function exists. In an equipment business the answer is unambiguous: support exists to sell the next machine and to keep the dealer selling your line.

18. Post-Service Repurchase Rate, and Dealer Renewal

What it is. The share of customers who buy again after a significant service event, against those who had no such event; and the share of dealers who keep your line, measured by share of counter rather than by contract status.

Formula. $\text{Repurchase} = (\text{accounts with a major service event in period T that ordered new equipment by T plus N}) / (\text{accounts with a major service event in T})$, times 100, compared against accounts with no such event. $\text{Share of Counter} = (\text{your revenue at that dealer}) / (\text{his total category revenue})$.

Benchmark. No public benchmark exists for either, and comparison to another company would be meaningless because purchase cycles differ by product. The economics of retention are sourceable: Reichheld and Sasser, in "Zero Defections: Quality Comes to Services" (Harvard Business Review, September/October 1990), reported that cutting the defection rate by 5% raised profits by 85% in one bank's

branch system, 50% in an insurance brokerage and 30% in an auto service company. For the dealer side, the NAEDA Cost of Doing Business Study, 2024 Edition put parts and service absorption at 40.02% for North American farm equipment dealers in the US\$200 million to US\$600 million band, with parts gross margin of 33.54% and inventory turns of 2.82 against a stated target above 3. That is the economics your dealer is managing, and it is more useful to know than any retention average.

How it lies to you. Long purchase cycles create attribution problems: a customer whose breakdown was in 2024 and whose replacement came in 2027 gets counted under 2027. Dealer retention is worse, a lagging indicator with a step function shape. A dealer stops promoting your line, stops stocking your parts and stops training on your equipment, all while remaining formally retained, and by the time the contract is not renewed the revenue left two years earlier.

What to do with it. Run the cohort comparison once a year: repurchase rate for accounts whose machine-down event was resolved inside your commitment, against those resolved outside it. That converts your support budget from a cost line into an investment case. For dealers, track share of counter and monthly parts trend, not contract status.

What to measure when support runs through the dealer and not through you

This is the most common situation in channel businesses and the least written about. The customer calls the dealer, the dealer's technician goes out, you never see the event, and your support dashboard is clean and mostly fictional.

19. Dealer-Handled Resolution Rate

What it is. The share of field events the dealer network resolves without escalating to the factory.

Formula. (Field events closed by the dealer) divided by (total field events on your equipment), times 100. The denominator is the hard part and where the discipline lies.

Benchmark. No public benchmark exists. Aberdeen's 70% first-time fix threshold is the closest defensible reference.

How it lies to you. A high dealer-handled rate is presented as network strength. It is equally consistent with a network that stopped escalating because your factory support is slow, or with dealers absorbing repeat failures and never telling you about a systemic defect. The number cannot distinguish a capable network from a resigned one.

What to do with it. Read it against escalation rate and warranty claim rate together. High resolution, low escalations and rising claims means that dealer is carrying a product problem for you. High resolution with low claims means he is either excellent or invisible, and the difference is worth a phone call.

The five things you can measure through the dealer

- Parts orders per unit in service, by dealer. A dealer well below the network median is sourcing elsewhere, which means either the aftermarket took your consumables or your fill rate pushed him away. This is your most reliable proxy for service activity you cannot see.

- Warranty claim rate and mix per dealer. Claims are self-reported and biased, but the variance between dealers is informative even when the level is not. A dealer with anomalously low claims on an aging population is absorbing cost off-warranty, and he will remember.
- Warranty claim cycle time, submission to payment. Entirely within your control and the strongest single driver of how a dealer feels about carrying your line. Measure the 90th percentile, not the mean.
- Technician certification coverage. Certified technicians per unit in service in that territory. Falling coverage on a growing installed base is a first-time fix collapse that has not happened yet.
- Escalation quality, not volume. When a dealer escalates, measure how long you take to give him an answer he can act on. Dealers stop escalating when escalating does not help, and once they stop you lose your only window into field failures.

The one thing you must buy, not infer

Those five are proxies. Direct measurement requires the end customer's voice, and in a dealer channel you acquire it deliberately: a short manufacturer-branded survey after a service event, run with the dealer's agreement and with results shared back to him.

Most manufacturers avoid this because it looks like going around the dealer. Structure it so it is not: he sees his own results and the network median, nothing about named competitors. He gets benchmarking of his service operation that he cannot build himself, and you get the one measurement your proxies cannot give you: whether the end customer would buy your equipment again.

A working scorecard

If you are starting from nothing, build eight metrics, and make them the ones that constrain each other.

Metric	Reported as	Paired with	Why the pair
In-window equipment uptime	Percent, inside the customer's critical season	Downtime hours by cause	The percentage says how bad, the causes say where
First-time fix rate	Percent, with cause codes on failures	Time to resolution, customer clock	Each blocks the other's gaming
Time to part on the ground	90th percentile, door to machine	Fill rate on machine-down parts only	Aggregate fill rate hides the parts that matter
Reopen rate	Percent, linked by serial and failure code	First contact resolution	The integrity check on every closure metric
Warranty claims rate	Percent of revenue, by production cohort	Warranty claim cycle time to dealer	Cost control that damages the channel is not control
CES after machine-down events	Top-two-box, high severity only	Survey response rate	A score without a response rate is not a measurement
Parts orders per unit by dealer	Index against network median	Technician certification coverage	The two leading indicators of a dealer drifting away
Post-service repurchase rate	Cohort comparison, annual	Dealer share of counter	The only two numbers that convert support into revenue

The discipline underneath all of it

Never reward a single metric in isolation. Every KPI here can be moved without improving anything, and a team under pressure finds the shortcut faster than you find the fraud. Pay on pairs that constrain each other.

Measure the customer's clock, not yours. Nearly every gap between a good support dashboard and an unhappy customer has one source: your timers start later and pause more often than his. He is counting from the moment the machine stopped, and he is not pausing for anything.

Go find the silent accounts. Support metrics are built entirely from customers who contacted you. The ones who did not are not satisfied, they are gone, and they are invisible in every number on your dashboard. In an equipment business they are also the largest pool of recoverable parts revenue you have. Nobody finds them in the data. You find them by calling.

Andre Magrini is a fractional Chief Revenue Officer based in Chicago, working with equipment manufacturers, agtech companies and industrial businesses on revenue architecture, channel performance and US market entry.