Customer Success and Retention KPIs

Seventeen retention metrics, which ones survive outside subscription, and what replaces the ones that do not

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

Every book about customer success assumes you sell a subscription. It assumes monthly recurring revenue, a product that emits telemetry, a health score built from login counts, and a renewal date on a calendar. That world is real, well documented, and not the world most revenue leaders live in.

I have spent most of my career on the other side of that line: equipment sold through dealers and distributors, OEM programs, feedlots. When I took the Brazil operation of Ag Growth International from roughly US$35M to more than US$150M in three years, none of it was subscription revenue. Retention was still the whole game, under a different name: the next purchase, the parts in between, and the dealer who keeps carrying your line.

This book covers both worlds. Seventeen KPIs, each with what it is, the formula, a sourced benchmark, how that KPI lies to you, and what to do with it. Where no credible benchmark exists, I say so instead of inventing one.

Two retention economies

Contractual recurring. Revenue renews on a date. Non-renewal is an observable event, so churn is countable.

Non-contractual repeat. The customer never cancels anything. They simply do not come back, and you find out eighteen months later when a competitor's machine is sitting in their yard. Departure is inferred from silence.

Almost every metric in the customer success canon was built for the first economy. Some transfer with modification. Some do not transfer at all, and pretending otherwise produces a green dashboard over a bleeding business.

Part one: the twelve KPIs of contractual retention

1. Logo churn rate

What it is. The percentage of customer accounts that stop being customers in a period, counted as accounts, not dollars.

The formula. Logo churn = customers lost in period / customers at start of period.

Benchmark. Recurly Research, Churn Rate Benchmarks (July 2026 data, Recurly's subscription network): total churn 3.60%, of which 2.34% voluntary and 1.25% involuntary. SaaS median 3.22%, top quartile 1.78%. Recurly's own caveat is right: a blended average mixing premium B2B with low-price consumer products describes nothing.

How this KPI lies to you. It treats a US$2,000 account and a US$400,000 account as one unit each. Lose two anchor accounts and keep forty small ones and the number looks fine while revenue collapses.

What to do with it. Never report it alone. Put it beside revenue churn, segmented by revenue band.

2. Revenue churn rate

What it is. Recurring revenue lost to cancellations and downgrades, before expansion is added back.

The formula. Gross revenue churn = (cancellation losses + downgrade losses) / recurring revenue at start of period.

Benchmark. Benchmarkit's 2025 SaaS Performance Metrics Benchmark Report (2024 data, private B2B software) reports median gross revenue retention of 88%, implying median gross revenue churn near 12% annually, down from 90% retention three years earlier.

How this KPI lies to you. Downgrades usually go uncounted. A customer cutting from 400 seats to 90 has churned 77% of their revenue while staying a happy logo on your customer slide. And it lands on the renewal date, not the decision date.

What to do with it. Split it into cancellation and contraction churn. Cancellation is an outcome failure; contraction is a value-per-unit failure or a budget event.

3. Gross revenue retention

What it is. The share of the starting revenue base still held at period end, counting only losses. Expansion is excluded, so GRR never exceeds 100%.

The formula. GRR = (starting recurring revenue - cancellations - downgrades) / starting recurring revenue.

Benchmark. Benchmarkit 2025: median 88%, down from 90% three years earlier. SaaS Capital's 2025 SaaS Retention Benchmarks for Private B2B Companies (firms above US$1M ARR) finds gross retention rises with average contract value.

How this KPI lies to you. GRR lies less than any other retention metric, which is why companies bury it. The distortion is denominator gaming: excluding accounts that were "never really our ICP" turns a bad quarter into a good one without changing a single relationship.

What to do with it. Define the denominator once, in writing, and never change it. GRR is the right metric for a customer success compensation plan: the team cannot inflate it through upsell.

4. Net revenue retention

What it is. The same starting cohort a year later, with expansion, upsell and price increases added back. Above 100% the base grows without a single new customer.

The formula. NRR = (starting recurring revenue + expansion - downgrades - cancellations) / starting recurring revenue.

Benchmark. SaaS Capital 2025: median NRR of 102% for companies with average contract values of US$25,000 to US$50,000, top quartile 111%, bottom quartile 97%. Benchmarkit 2025 puts the overall median at 101%. Both sit far below the 120% that circulates in pitch decks, which describes outliers.

How this KPI lies to you. NRR is a weighted average dominated by your largest accounts. If your top ten expand aggressively while your long tail churns at 30% a year, NRR reads 105% while the mid-market motion dies. Second lie: a price increase, a contractual escalator and genuine usage growth count identically. NRR cannot separate revenue you earned from revenue you invoiced.

What to do with it. Report it three ways: overall, excluding the top decile, and split into price-driven versus volume-driven expansion. Always publish GRR beside it. The gap is the size of the story NRR is hiding.

5. Expansion rate and upsell contribution

What it is. How much of your new revenue came from customers you already had.

The formula. Expansion contribution = expansion revenue / (new logo revenue + expansion revenue).

Benchmark. Benchmarkit 2025: expansion contributes a median of 40% of new ARR, up five points year over year, and 67% above US$100M ARR. Few companies apply to expansion the rigour they apply to new logo revenue.

How this KPI lies to you. Expansion is the easiest number in the business to manufacture. Raise list price 8% across the base and you have booked expansion while customers are unhappy. High expansion with falling gross retention is a base being harvested.

What to do with it. Tag every expansion dollar with a cause at booking: new users, new module, new site, usage growth, price. Price-only expansion above a quarter of the total is a warning, not a win.

6. Renewal rate

What it is. Of the contracts up for renewal, the share that renewed. The denominator is only accounts with a renewal event.

The formula. Gross renewal rate by value = value of contracts renewed / value of contracts up for renewal in the period.

Benchmark. None I would put my name behind, and the reason matters. Published studies (SaaS Capital, Benchmarkit) measure against the full base, not the renewal cohort, so their figures are not comparable. Nearly every renewal benchmark in circulation is a retention figure relabelled.

How this KPI lies to you. It is the most timing-manipulable metric in the set. A customer who was leaving gets a two-month extension and the loss moves into next quarter. And it is blind to the negotiation: a renewal at 60% of prior value still counts.

What to do with it. Measure by value, not count. Track "renewed at or above prior value", and track extensions separately. A rising extension count is the cleanest leading indicator of a bad quarter ahead.

7. Customer health score

What it is. A composite index combining usage, support, relationship and commercial signals into one number meant to predict renewal.

The formula. There is no standard formula. It is always a weighted sum the company defines, with weights set by judgment.

Benchmark. No benchmark exists and none can. Every company defines inputs and weights differently, so a 72 here and a 72 there describe nothing in common. The most quoted accuracy figure is that 73% of customer success leaders say their health score does not reliably predict churn, attributed to ChurnZero's 2025 Customer Success Leadership Study of roughly 800 post-sales leaders. I could not open the primary study, so treat it as directional.

How this KPI lies to you. A health score built mainly on product usage measures the wrong human being. Usage is generated by users. Renewals are signed by economic buyers. I have watched accounts with heavy daily engagement cancel because a new CFO consolidated vendors, and the usage graph gave zero warning, because nothing in it touched the CFO.

What to do with it. Split it into a churn risk score and an expansion readiness score, and include at least one variable about the economic buyer. Then each quarter, check what health score each churned account carried ninety days earlier. If it was green, the score is decoration.

8. Time to first value

What it is. Elapsed time between signature and the first outcome the customer bought the product for. Not first login. Not go-live.

The formula. Time to first value = defined value milestone date minus contract start date, reported as a median, because the long tail destroys the average.

Benchmark. None I would publish. Figures in circulation come from vendor blogs with undisclosed sample and value definitions, and they disagree by orders of magnitude. Because every company defines value differently, this metric is not comparable across companies.

How this KPI lies to you. Most teams measure time to launch and call it time to value, because launch has a date and value does not. Launch is when your work finished. Value is when their work got better. The gap is where churn is born.

What to do with it. Write down, per segment, the observable event that constitutes first value, and get the customer to agree to it during onboarding. Report the median and the 90th percentile, which is where your churn lives.

9. Product adoption rate

What it is. The share of licensed capacity actually used, by seats, features, or depth of workflow.

The formula. Seat adoption = active users / licensed users. Feature adoption = accounts using a feature / accounts entitled to it.

Benchmark. No credible cross-company benchmark exists, for the same reason as health score: "active" runs from one login in ninety days to daily use of a core workflow. Any number quoted without its activity definition is noise.

How this KPI lies to you. High adoption is not high value: heavy usage of a workaround looks identical to heavy usage of a solution. Adoption also flatters you when the buyer over-bought: 90% adoption of 50 seats when the customer paid for 200 is a contraction event waiting for the renewal date, invisible because most teams count deployed rather than purchased seats.

What to do with it. Always calculate against contracted capacity. Pair it with breadth: the number of distinct departments or sites using the product predicts renewal better than intensity inside one team.

10. Economic buyer engagement versus user engagement

What it is. Two measures kept separate on purpose. User engagement is what daily users do. Economic buyer engagement is whether the budget holder has had an outcome-level conversation with you inside a defined window.

The formula. Economic buyer coverage = accounts with a documented economic buyer contact in the last 90 days / all accounts.

Benchmark. Gartner, surveying 632 B2B buyers in August and September 2024, found buying groups of five to sixteen people across as many as four functions, and reported in May 2025 that 74% of buying teams show unhealthy conflict while deciding. If five to sixteen people decide, knowing two of them is not coverage.

How this KPI lies to you. Fifteen names in the CRM, fourteen cc'd on a status email, is not multithreading. What matters is whether you have discussed their business outcome with someone who could cancel you.

What to do with it. Set a hard rule: no account above a revenue threshold goes two quarters without a documented economic buyer conversation. Track champion departures and trigger a playbook the same week. Dramatic post-departure churn figures circulate, but they come from vendor marketing without methodology.

11. QBR coverage and QBR quality

What it is. Whether structured business reviews happen where they are needed, and whether anything measurable comes out.

The formula. QBR coverage = accounts with a completed review in the last two quarters / accounts in the tier. QBR yield = reviews producing a dated commitment / reviews held.

Benchmark. No published benchmark for QBR frequency or effectiveness that I would rely on. What exists is vendor guidance, not measurement.

How this KPI lies to you. Coverage is a pure activity metric, and activity metrics reward the meeting rather than the outcome. A review where you show a usage dashboard to three users, none controlling budget, counts as full coverage and is worth nothing.

What to do with it. Measure four things inside the review instead of the review itself: was the economic buyer present, was an outcome quantified in the customer's own numbers, and did they commit to a dated action.

12. Advocacy: NPS, referral rate and earned growth

What it is. Whether customers will put their own name behind you. NPS measures intent. Referral and earned growth measure behaviour.

The formula. NPS = % promoters (9 to 10) minus % detractors (0 to 6). Earned growth rate, per Fred Reichheld, Darci Darnell and Maureen Burns in "Net Promoter 3.0" (Harvard Business Review, November and December 2021), = net revenue retention plus earned new customers, minus 100%.

Benchmark. Survicate's 2025 NPS Benchmarks (599 companies, 5.4 million responses) reports a median NPS of 38 for B2B against 49 for B2C. MeasuringU's Business Software UX and NPS Benchmarks (2025) reports an average of -5% across 23 business software products. Same metric, same industry, 46 points apart. On behaviour, Reichheld, Jamie Cleghorn and Wojtek Kokoszka in Harvard Business Review (September and October 2026), analysing more than 10 million consumers, report that about 20% of new customers arrive through referral. Nielsen's Trust in Advertising Study (2021, 40,000 consumers) found 88% trust recommendations from people they know.

How this KPI lies to you. NPS lies through sampling: the customers who answer are the ones with a relationship, and those are the ones who were not leaving. Keiningham, Aksoy, Cooil and Buoye ("The Wallet Allocation Rule", Harvard Business Review, October 2011) found no managerially relevant correlation between satisfaction or NPS and share of wallet. A customer can rate you a 9 and give you 30% of their spend, because they rate your competitor a 9 too.

What to do with it. Stop treating NPS as a score to raise and treat it as a list of names to call. Then add a behavioural metric on top. Earned growth is the best candidate: it is built from accounting outputs, not survey inputs.

Retention when there is no subscription

Now the part nobody writes.

If you sell a machine every seven years, or a capital system with a decade of life, or product through a dealer network that owns the end customer, most of Part One does not apply. Monthly churn does not exist because there is nothing to cancel. Usage health scores do not apply: where telemetry exists, heavy use signals a productive customer, not a loyal one. Renewal rate does not exist because there is no renewal date. Three layers replace them.

Layer one: the transaction cycle. The customer buys, goes quiet for years, then buys again or does not. Retention means winning the next transaction, measured over the cycle rather than the quarter.

Layer two: the revenue between transactions. Parts, service, consumables, wear items, upgrades, training. This is where the relationship is maintained and usually where the profit lives. A customer buying parts elsewhere has already partly left, and the machine in their yard is telling you nothing.

Layer three: the channel. Selling through dealers stacks two customers on top of each other. You can keep the end user and lose the dealer, and then you lose the end user next cycle anyway.

13. Repurchase rate and interval between purchases

What it is. The share of customers who buy again within the expected replacement cycle, and the time between purchases. The non-contractual equivalent of renewal rate.

The formula. Repurchase rate = customers who purchased again within the cycle window / customers eligible to repurchase in that window. Eligibility, not the calendar, sets the denominator: a customer who bought a combine last year is not eligible this year and must not drag the rate down.

Benchmark. The closest usable proxy in agricultural equipment is brand loyalty. Ag Equipment Intelligence's brand loyalty study found 75% of farmers describing themselves as brand loyal in 2017 and 63% in 2020, with its 2024 work finding loyalty had not returned to the 2017 peak. Self-declared loyalty is intent, so treat it as a ceiling. On the cycle, Iowa State University Extension's Ag Decision Maker (A3-30, Replacement Strategies for Farm Machinery) puts the optimal replacement point at the year annualised total cost begins to rise, around year nine in its example. There is no industry replacement-cycle number to borrow: off-road equipment is not registered like vehicles, so as Farm Equipment notes, fleet age cannot be stated with confidence.

How this KPI lies to you. Measured over any window shorter than the replacement cycle it is noise, moving on commodity prices, interest rates, depreciation rules, a drought. It is dominated by whichever cohort sits inside its window this year, which makes year-over-year comparison a trap.

What to do with it. Measure by purchase cohort, not calendar year. Track the interval trend rather than the rate: an interval stretching from 6.5 to 8 years is a slow-motion revenue loss no annual number shows. Flag every repurchase that needed an out-of-policy discount. That is a retention failure wearing a sale's clothing.

14. Parts and service attach rate

What it is. Attach rate is the share of your installed base buying parts, service or consumables from you rather than a third party. Absorption is the dealer-level sibling: operating expenses covered by parts and service gross profit.

The formula. Parts attach rate = units generating parts revenue / units in service. Absorption rate = (parts gross profit + service gross profit) / total operating expenses.

Benchmark. The Western Equipment Dealers Association's Cost of Doing Business Study (2020 data, reported by Farm Equipment) found an average absorption rate per location of 73.2%, ranging from 52.7% at the smallest dealers to 76.5% at the largest, with parts at 16.7% of store sales carrying 35.3% gross margin. The United Equipment Dealers Association targets 80% or better for farm and outdoor power equipment dealers. McKinsey reports an average EBIT margin of 25% for aftermarket services against 10% for new equipment across 30 industries; Deloitte puts aftermarket operating margin at roughly 2.5 times that of new equipment. On attach rate itself there is no published cross-industry benchmark.

How this KPI lies to you. Attach measured in revenue rather than units hides the shape of the problem. Ten large customers buying heavily produce a healthy revenue attach while 60% of your installed base has not bought a part in three years, and that 60% buys the competitor's machine next cycle. Warranty work is obligation, not loyalty.

What to do with it. Measure unit attach on out-of-warranty units only. Build the list of silent units and treat it as a churn list. On the dealer side, absorption below 80% is structural fragility: those fixed costs ride on whole goods, and a downturn forces decisions that damage your brand.

15. Dealer and distributor retention

What it is. Whether the partners carrying your line keep carrying it, keep it at the front of the yard, and keep training people on it.

The formula. Dealer revenue retention = this year's revenue from last year's dealers / last year's revenue. Alongside it, share of shelf: your line as a percentage of the dealer's category revenue.

Benchmark. There is no benchmark for dealer retention rate. What exists is the consolidation trend that sets the denominator. Ag Equipment Intelligence and Currie Management Consultants, tracking farm equipment dealer ownership in the US and Canada since 2009, report that "big dealers" (five or more ag locations) peaked at 214 in 2022 and fell to 199 by 2026, a fourth consecutive annual decline. John Deere led at 77 in 2026, down 26% since 2012.

How this KPI lies to you. In a consolidating channel, count-based retention falls every year for reasons unrelated to your performance, because two of your dealers merged into one. And a dealer who has not formally dropped you can be functionally gone: still holding the agreement, still ordering occasionally, quietly no longer recommending you.

What to do with it. Retire count-based retention and measure revenue retention with mergers reconciled into one entity. Instrument three leading indicators: parts stocking depth, current technician certifications, and share of the dealer's category revenue. A dealer falling on the first two has already left.

16. Share of wallet at the customer

What it is. Of everything this customer spends in your category, the share that comes to you. The only retention metric that counts the competitor inside the same account.

The formula. Share of wallet = your revenue from the customer / their total category spend. Since total spend is rarely observable, the practical instrument is the Wallet Allocation Rule: (1 minus (rank / brands used)) multiplied by (2 / (brands used plus 1)).

Benchmark. The source rather than a number: Keiningham, Aksoy, Cooil and Buoye, "The Wallet Allocation Rule", Harvard Business Review, October 2011. What predicts share is rank and the number of brands the customer uses, not satisfaction. No universal benchmark exists: the number is a function of category structure.

How this KPI lies to you. Every proxy is estimated, usually by the account manager compensated on the account, which makes it optimistic. Estimates from your own sales data are worthless: that data only sees what the customer bought from you. Second failure: total share at a multi-site customer can read 55% while one division is at 100% and three are at zero.

What to do with it. Get rank from the customer directly, as two questions in the review: of the brands you use here, where do we sit, and how many do you use? Track it at site or division level, never at parent level.

17. Contactable installed base coverage

What it is. The share of units you have sold that you can reach: you know who owns the machine today and have a named person who will take your call.

The formula. Contactable coverage = units with a verified owner and a contact touched in the last 12 months / units still in service.

Benchmark. No published benchmark exists, and the reason is structural: in channel-sold equipment the manufacturer frequently does not own the end customer record, and used machines change hands without anyone telling the manufacturer.

How this KPI lies to you. Installed base is normally reported as units shipped, a cumulative number that only rises and therefore cannot report a problem. It counts machines that were scrapped, resold twice, or whose owner you could not name under oath.

What to do with it. Separate units shipped, units in service and units contactable, and report the third. In a channel business, write visibility into the end customer record into the dealer agreement, because the alternative is a business whose customer list belongs to someone else.

The exit signals that show up before churn

Churn is an event. Leaving is a process, observable months earlier if you know what to watch. The signals differ by business type, and the wrong set is worse than none.

Business type	Leading signals, earliest first	Warning window
Subscription software	Economic buyer stops attending reviews. Champion changes role or employer. A new executive arrives above your buyer. Support tickets shift from "how do I" to "why does this". Breadth of use narrows to one team while total usage holds.	Two to four quarters

Business type	Leading signals, earliest first	Warning window
Contract services and retainers	Meeting cadence slips and they do not reschedule. Scope questions become budget questions. The client starts doing internally a piece of what you do.	One to three quarters
Capital equipment, direct	Out-of-warranty parts orders stop or move to third parties. Service calls decline while machine hours hold. The customer stops attending training. A competitor demo unit appears on site.	One to three years, before the replacement decision
Dealer and distributor channel	Parts stocking depth falls. Technician certifications lapse. Your line moves to the back of the lot. The dealer takes a competing line "for a segment we do not cover". Ownership change at the dealership.	Two to six quarters
OEM and integrator relationships	Your component stops appearing in new platform designs while current programs continue. Engineering contact drops out of technical reviews. Forecast horizon shortens from annual to quarterly.	One to two product cycles

Two rules cut across all five rows. First, the earliest signal is relational, not transactional. The person stops before the money stops. Second, silence is data: a customer who has not been in touch is not a satisfied customer, they are an unmeasured one.

What transfers and what does not

KPI	Outside subscription	What replaces or modifies it
Logo churn (monthly)	Does not transfer	Repurchase rate over the replacement cycle
Revenue churn (monthly)	Does not transfer	Aftermarket revenue per unit, by cohort
Gross revenue retention	Does not transfer	Same-customer revenue year over year, cycle-adjusted
Net revenue retention	Does not transfer as stated	Share of wallet plus attach rate, measured per site
Renewal rate	Does not transfer	Repurchase rate on the eligible cohort only
Product adoption rate	Does not transfer	Utilisation and uptime of the unit
Usage-based health score	Does not transfer	Attach rate, contactability, relationship depth
Time to first value	Transfers, modified	Commissioning to first productive season
Economic buyer engagement	Transfers directly	Same discipline, longer intervals
QBR discipline	Transfers, modified	Annual review timed to their planning season

KPI	Outside subscription	What replaces or modifies it
Advocacy and referral	Transfers directly	Often stronger in equipment than in software
Expansion revenue	Transfers, redefined	Attach, upgrades, additional units and sites

Building the scoreboard

For a contractual recurring business, the minimum honest set is five: gross revenue retention, net revenue retention, logo retention by revenue band, expansion split by cause, and economic buyer coverage.

For a non-contractual business, also five: repurchase rate on the eligible cohort, purchase interval trend, out-of-warranty unit attach rate, dealer revenue retention with mergers reconciled, and contactable installed base coverage. Add share of wallet by rank when your account managers will ask those two questions.

One discipline applies to both. Once a quarter, take everything you lost and check what your metrics said ninety days earlier. If the dashboard was green, the dashboard is decoration. That test costs an afternoon, and it is the only item here that tells you whether the rest works.

Andre Magrini is a fractional Chief Revenue Officer in Chicago. He led the Brazil operation of Ag Growth International from roughly US$35M to more than US$150M in three years, selling capital equipment through dealer and distributor channels, to OEMs and to feedlots.