

# NRR Leakage Audit

Nine dimensions of post-sale discipline. The audit that tells you where Net Revenue Retention is leaking before your next hire — and usually why "customer success is working" is not the same as "expansion is working."

## WHEN TO USE

Score across all 9 dimensions with the post-sale owner (VP CS, VP Rev, or whoever actually gets NRR in their comp). Use the shared 1-5 legend. Total the score. Read the band. The common diagnostic mistake: assuming strong gross retention means expansion is fine. It usually means churn is suppressed and expansion is invisible.

**Preamble.** Most GTM teams spend ~80% of budget on pipeline and treat post-sale as a cost center. The math on that allocation stops making sense the moment you run the NRR numbers: a team going from 104% to 118% NRR creates more compounding value than a 30% pipeline increase — at a fraction of the cost. This audit maps nine post-sale disciplines against that reality.

## Source lineage:

- Rackham SPIN — discovery structure (dimension 1)
- Christensen/JTBD — expansion is a new hiring event, not a cross-sell (dimensions 2, 4)
- Deming 85/15 — churn variation lives in the system, not the CSM (dimensions 3, 4)
- Kahneman & Tversky prospect theory — QBR framing as loss prevention vs. gain promise (dimension 5)
- Fisher & Ury — principled negotiation applied to expansion conversations (dimension 6)

## Scoring legend (1-5). Every dimension uses the same 1-5 scale:

- 1 — Absent:** You do not do this. Either you have not built it, you built it and retired it, or the discipline does not exist in the org.
- 2 — Ad-hoc:** Someone does this, sometimes, when they remember. No cadence, no owner, no artifact. Cannot be audited.
- 3 — Documented:** An owner exists. A cadence exists. An artifact exists. Not yet measured; not yet feeding other layers.
- 4 — Measured:** Outputs are tracked against a baseline. Reviewed at least quarterly. Decisions are made from the numbers, not vibes.
- 5 — Compounding:** Outputs feed back into upstream layers. Every cycle makes the next cycle sharper. The discipline survives the person who built it.

Total possible: 45. Your band determines where to look first.

Score range	Band	Where to look first
≤ 18	Foundation missing	NRR is likely <100% and you cannot explain why. Stop hiring in CS. Instrument cohort-NRR by PSP (dimension 1) and a 5-signal churn-risk system (dimension 4) before anything else. Most leaks at this level are visibility problems, not execution problems.
19 – 29	Leaking between layers	NRR is somewhere in the 95-105% band and fragile. The typical leak is QBR structure (dimension 5) and account planning (dimension 3). Rebuild QBRs as EVP-delivery sessions this quarter; expansion bookings per QBR will 2-3x within one cycle, which changes the math entirely.
30 – 38	Working, not compounding	NRR is 105-115% but not yet compounding. The gap is usually multi-threading (dimension 6) and outcome-level metrics (dimension 7). These are the two disciplines that separate a "working" CS motion from a compounding one — threading prevents churn-by-attrition, outcome metrics unlock expansion conversations.

39 – 45	Compounding	NRR is 115%+ and the motion is compounding. The risk now is success-plan operator assignment (dimension 9) breaking at scale — as the book grows, unowned accounts appear silently. Quarterly book audits catch this early; without them, NRR declines 6-12 months before anyone notices.
------------	-------------	---

---

 DECISION CRITERIA
 

---

### 1. Segment cohort NRR tracking (by PSP, not company size)

*NRR reported by ACV tier or company size masks the interesting variance. Cohort NRR by PSP is where you discover that one PSP is compounding at 130% while another is leaking at 92% — which tells you which PSPs to double-down on and which to retire.*

**Score 1 (low):** NRR tracked as a single blended number, or at most by ACV tier. No PSP-cohort view exists.

**Score 5 (high):** NRR tracked per PSP cohort, quarterly. Variance drives hiring, budget, and retention decisions. Low-NRR PSPs are retired or redesigned.

**Recommendation:** Tag every customer with the PSP they came in on. Build a cohort-NRR report this quarter. Run it at the next leadership review; the conversation about which PSPs to keep will change.

### 2. Onboarding-to-activation cycle time

*The time from close to first measurable value. Christensen's point applies: the customer hired you for a specific job. Activation is proving the job got done. Long onboarding cycles destroy expansion because the first EVP promise never gets validated.*

**Score 1 (low):** Activation is undefined. Time-to-activation is unmeasured. Onboarding = a 90-day welcome sequence.

**Score 5 (high):** Activation defined per PSP as a measurable milestone proving one EVP promise. Time-to-activation tracked and benchmarked. <45 days for most segments.

**Recommendation:** Define activation per PSP this week: one measurable milestone that proves the pre-sale EVP promise. Measure time-to-activation on every account. Any account past 60 days with no activation is at-risk, regardless of NPS.

### 3. Expansion account plan quality

*Does every account have a named next-EVP hypothesis? Expansion is not cross-sell; it is the customer hiring the next job at the moment the previous one was completed. Without a named next-EVP, expansion becomes reactive (they churn before we pitch).*

**Score 1 (low):** No account plans. Or plans are CRM notes with "expansion opportunity: TBD" populated manually once a quarter.

**Score 5 (high):** Every account has a named next-EVP, tied to a PSP signal. Account plans reviewed quarterly. Expansion forecasts tie to specific signals, not CSM vibes.

**Recommendation:** Build a 1-page account plan template: current-EVP proof, next-EVP hypothesis, target trigger event, projected ACV impact. Require every >\$50K account to have one by end of quarter.

### 4. Churn-risk signal system

*Observable leading indicators of churn, not lagging ones. Usage drops, seat-removal, champion departures are all leading. NPS and "CSM gut feel" are lagging — by the time they fire, the deal is already in the save motion, not the retention motion.*

**Score 1 (low):** Churn detected via CSM intuition or NPS scores. No observable signal system. Saves happen at renewal, not before.

**Score 5 (high):** 5-8 leading-indicator signals automated, firing alerts 90+ days before renewal. Playbook per signal. Success rate on signals tracked.

**Recommendation:** Pick 5 leading-indicator signals this week (usage drop, login decay, champion change, seat removal, support-ticket volume). Instrument alerts. Build a playbook per signal. The 90-day lead time on an at-risk account is the difference between retention and save-motion.

### 5. QBR structure + discipline

*QBRs as EVP delivery sessions, not status updates. The first 25 minutes: proof the promised EVP was delivered, quantified. The second 35 minutes: structured surfacing of the next EVP. Status-update QBRs generate nothing; EVP-delivery QBRs compound expansion.*

Score 5 (high): QBRs run to a fixed agenda: 20 min EVI delivery prep + 30 min next EVI packaging + 20 min status updates. Expansion bookings per QBR tracked.

**Recommendation:** Rewrite the QBR agenda template this week: proof (25 min), next need (35 min), zero status. Train CSMs on the new agenda. Expansion per QBR is measurable within one cycle; you'll see 2-3x within a quarter if the new format sticks.

## 6. Multi-threaded account coverage

*How many named relationships per account? Single-threaded accounts are churn-by-default when the champion leaves — and every champion leaves eventually. Multi-threading is not "LinkedIn-connecting with 10 people"; it is having spoken to 3+ named stakeholders in the last 90 days.*

**Score 1 (low):** Most accounts single-threaded. One champion per account. No threading strategy.

**Score 5 (high):** 3+ active relationships on every >\$50K account. Threading tracked per account. Stakeholder changes trigger re-threading playbook within 14 days.

**Recommendation:** Run a threading audit across top-20 accounts this week. Any account with <3 named relationships is at elevated risk regardless of health score. Put a 14-day re-threading SLA on any stakeholder departure.

## 7. Adoption metric depth

*Are adoption metrics feature-level or outcome-level? Feature-level metrics ("% of seats using search") measure the tool; outcome-level metrics ("average meetings booked per rep using the tool") measure the value. Outcome-level is the only metric that supports expansion conversations.*

**Score 1 (low):** Adoption = logins or seat count. No outcome-level metrics instrumented. Cannot quantify value delivered.

**Score 5 (high):** 3-5 outcome-level metrics per PSP, instrumented in the product or via QBR capture. Benchmarks published per PSP. Customers see their own metrics vs. the benchmark.

**Recommendation:** Pick 3 outcome-level metrics per PSP this quarter. Instrument them. Benchmark them. Customers who see their own outcomes vs. a peer benchmark expand at ~2x the rate of customers who only see raw adoption numbers.

## 8. CS-to-Sales handoff quality

*When expansion opportunity surfaces, how is it routed back to Sales? "Sales does expansion" vs. "CS does expansion" is the wrong debate. The right question: what is the structured handoff when CS surfaces a qualified expansion signal?*

**Score 1 (low):** No defined handoff. CS either does expansion badly or throws it over the wall with no context.

**Score 5 (high):** Structured handoff: CS qualifies + packages the signal; Sales carries the commercial conversation. SLA on handoff response. Handoff-to-close conversion tracked.

**Recommendation:** Define a CS→Sales expansion handoff template this quarter: PSP, signal, qualified budget range, next-EVP hypothesis, champion relationship. Track handoff-to-close conversion. The conversation about "who owns expansion" goes away once the handoff has a SLA.

## 9. Success plan operator assignment

*Is every active account assigned to a named operator accountable for the success plan? Deming's 85/15 rule: most NRR variation is system, not individual — but the system only works when the individual is named. Unowned accounts decay at ~3x the rate of owned accounts.*

**Score 1 (low):** Most accounts have no named CS owner. Or ownership rotates opaquely. Accounts fall through the cracks.

**Score 5 (high):** Every active account has a named CS owner. Ownership changes are documented and communicated to the customer within 14 days. Book-of-business reviewed quarterly.

**Recommendation:** Audit your book for unowned or recently-reassigned accounts. Any account without a named owner for 30+ days is a retention risk; rehome within a week. Communicate ownership changes to the customer proactively — silent changes destroy trust.