


Building Performance-Driving Compensation Plans

FEATURING

Worked examples, a 10-minute diagnostic, and a pre-launch stress test for AE, CSM, and SDR/BDR plans that drive performance.



In This Report

Most sales compensation plans clear finance's cost model and legal's fairness bar, then quietly fail to move the two things that we think matter most: **revenue quality** and **rep urgency**.

Our latest report shares:

- 1 Five principles behind comp plans that *do* drive performance
- 2 A 10-minute diagnostic to score your current AE or CS plans
- 3 Three worked examples (Account Executive, Customer Success Manager, and SDR/BDR) with payout math
- 4 A five-part stress test to run before any plan goes live.

Why Most Comp Plans Don't Drive Performance

According to [RepVue's Cloud Sales Index](#), average quota attainment across SaaS sales orgs in Q2 last year was 42.69%. That means most **reps closed less than half their quota that quarter**.

For reference, we think that a healthy attainment number is 80%.

When average attainment falls below ~50%, that's often a sign that you set the quota way higher than what your team can realistically produce.

So, when you think you have a team performance issue, consider pointing your finger at the comp plan and quota instead.



Where comp design most often fails

- Too many metrics diluting focus
- Misaligned incentives
- Pay mix that doesn't match how much control the role has over the outcome
- Quotas set top-down instead of from territory or account capacity
- Thresholds and accelerators that are missing or too weak to change behavior

These factors lead to a lack of trust across your selling teams, and without that trust and buy-in, your comp plan won't drive performance.

A 3-Point Diagnostic:

Score Your Plan in 10 Minutes

To help measure how effective your current comp plan is in driving performance, our proprietary AI revenue strategist, Atlas, created a scorecard that gives you a quick read on your AE or CS plan.

You'll score each category from 1 to 3.

Does the plan overpay for the revenue you want most? Ask:

- Does the plan pay differently for new logo versus expansion versus renewal?
- Does it reward the revenue attributes that matter most to economics: multi-year term, upfront cash, ICP or strategic segment, higher-margin product mix, self-sourced or outbound deals?
- If two deals have the same ARR but vary by quality, do they pay differently?

Healthy sign:

Reps can earn noticeably more on the company's "best" deal shape.

Warning sign:

Every booked dollar pays about the same, so reps optimize for speed, not quality.

Your
Score:

3 = clear differentiation tied to business value.
2 = some modifiers exist, but they're too small to change behavior.
1 = flat plan; all revenue is treated roughly the same.

Does the payout curve create urgency around quota? Ask:

- Is there a clear earnings difference between 90%, 100%, and 110% attainment?
- Do accelerators start at a point that motivates stretch behavior?
- Is the upside strong enough that top performers materially out-earn average performers?
- Is underperformance handled with a sensible decelerator or threshold?

Healthy sign:

The next dollar becomes more valuable at the right thresholds.

Warning sign:

The plan is too flat, so reps don't push late in the period or they defer deals across periods.

Your
Score:

3 = meaningful separation below, at, and above goal.
2 = accelerators exist, but not enough to change behavior.
1 = little economic difference across attainment bands.

Is the plan simple enough to be trusted? Ask:

- Can a rep explain how much they earn, for what, and when in under two minutes?
- Are payout timing and eligibility rules clear?
- Are crediting rules clean for splits, overlays, inbound versus outbound, and handoffs?
- Has the plan remained stable, or does it rely on frequent exceptions and midyear adjustments?

Healthy sign:

Reps trust the math and spend time selling, not auditing.

Warning sign:

Shadow spreadsheets, constant exceptions, and repeated payout disputes.

Your Score:

3 = simple, stable, explainable.
2 = understandable, but with recurring confusion.
1 = opaque, exception-heavy, low trust.

How to interpret the result

Score total	Readout
8-9	Strong plan architecture; likely performance-positive
6-7	Usable, but probably leaving performance on the table
3-5	Plan is likely creating drag, confusion, or poor revenue mix

A comp plan only does its job if it:



Improves revenue quality



Increases urgency around target



Is trusted enough that reps respond predictably



TIP



If you want a structured second opinion, Atlas can run this same diagnostic against your plan and flag the two changes with the biggest performance impact.

Five Principles of a Performance–Driving Plan

If your score comes back in a range that you're unhappy with, and you plan to adjust it, follow each principle below as you re-design.

1. Match pay mix to influence over the outcome

Current 2026 benchmarks:

- AE pay mix runs roughly 45/55 to 50/50 base to variable
- CSM is trending toward 70/30 as outcome-linked plans become the standard
- SDR typically runs 60/40 to 70/30 depending on inbound versus outbound motion.

The logic: the more a role can single-handedly move the outcome, the more of its pay should be at risk on that outcome.

2. Limit the plan to three measures/plan components

Protect your plan from doing too much by keeping plan components or measures to 3 (think: accelerators, bonuses, floors).

When you add a fourth measure or component, reps can no longer tell what actually earns them money, which quietly erodes the behavioral power of every measure in the plan.

3. Set a threshold/floor

A threshold is a set performance floor. If you fail to hit above the floor, you don't earn commission. Without one, a rep can collect a meaningful commission check from mediocre results... which ultimately removes the pressure a plan is supposed to create.

We've found a threshold is most appropriate for customer teams, like CSMs and account managers. For sales, a decelerator may be more appropriate (a decelerated commission rate until hitting a certain attainment band).

4. Build in leverage for over performance via tiered accelerators

A common structure pays 1.5x the standard rate at 100 to 125% of quota, and 2x at 125 to 150%.

Tiered acceleration is what makes hitting 140% of quota meaningfully more lucrative than hitting 100%, rather than just proportionally more.

You'll want to model and test these (can use Atlas for this) before implementing to avoid huge payouts.

5. Calibrate quota to territory and account reality

Avoid backing into a revenue target then setting your quota.

Instead, build from territory and account potential.

Why? A quota should be hard but achievable for a rep with a normal book that's rooted in math. If you focus only on the math without factoring total addressable market (TAM) and historical performance, you'll find your teams far below healthy attainment levels.

Three Worked Examples

Each of the three examples below considers the levers from the section above: **pay mix**, a **limited metric set**, a **real threshold**, a **tiered accelerator**, and **calibration** based on **how much control each role** actually has over its outcome.

That's what ties them together as **performance-driving plans** versus plausible pay structures.

1


**Mid-Market
Account Executive**

2

**Customer Success
Manager**

3

SDR/BDR



Mid-Market Account Executive

- \$140K OTE
- 50/50 split (\$70K base, \$70K variable)
- \$700K new ARR quota, a single metric (new ARR closed-won)
- a 10% commission rate
- a decelerator to 5% commission (half rate) below 50% attainment
- a 1.5x accelerator above 100% of quota.

ATTAINMENT	ARR CLOSED	VARIABLE PAYOUT	% OF TARGET VARIABLE	TOTAL COMP
40%	\$280,000	\$14,000	20%	\$84,000
65%	\$455,000	\$45,500	65%	\$115,500
100%	\$700,000	\$70,000	100%	\$140,000
140%	\$980,000	\$112,000	160%	\$182,000

It's a single clean metric because an AE has the most direct control over that outcome of any of the three roles, so there's no reason to dilute it with secondary measures.

The 50/50 pay mix puts real money at risk, which only makes sense because the AE can move the number.

The decelerator at 50% attainment does different work than a hard floor.

A rep sitting at 40% still earns something, \$14,000 instead of \$0, but at half the standard rate, so the gap between "underperforming" and "hitting the halfway mark" stays real without a cliff that pays nothing at all.

That matters behaviorally: **a true \$0 floor can tempt a rep who's clearly going to miss 50% into checking out entirely for the rest of the quarter**, since nothing they close moves their paycheck. A decelerator keeps every closed deal worth something, so there's still a reason to keep selling even in a rough stretch, just at a reduced rate until they clear the halfway point.

And the 1.5x accelerator above quota means the payout curve gets steeper than the performance curve once a rep clears 100%: going from 100% to 140% attainment is 40% more revenue, but it's 60% more commission.

That gap is what gives a rep who has already hit quota a reason to keep selling instead of pacing into next quarter. The plan is symmetric in effect if not in magnitude; rates soften below target and steepen above it, so the curve tracks performance in both directions instead of just penalizing the bottom.

Customer Success Manager

- \$130K OTE
- 70/30 split (\$91K base, \$39K variable)
- Managing roughly a \$3M ARR book

Two metrics: 60% weight on Gross Revenue Retention (target payout \$23.4K) and 40% weight on expansion or upsell ARR (target payout \$15.6K, on a \$150K expansion quota, at a 10.4% commission rate).

The retention component pays \$0 below 90% GRR, scales to full target payout at 95% GRR, and accelerates by 20% of target payout per point above target; the expansion component pays straight commission with a 1.5x accelerator above quota.

	GRR	EXPANSION ARR	RETENTION PAYOUT	EXPANSION PAYOUT	TOTAL VARIABLE	% OF TARGET	TOTAL COMP
BELOW THRESHOLD	90%	\$0	\$0	\$0	\$0	0%	\$91,000
ON-TARGET	95%	\$150,000	\$23,400	\$15,600	\$39,000	100%	\$130,000
OVERPERFORMER	99%	\$220,000	\$42,120	\$26,520	\$68,640	176%	\$159,640

The interesting design choice here **is splitting one metric into two** instead of collapsing everything into a flat renewal bonus, which is the old-school pattern this report argues against. Weighting retention 60% and expansion 40% says the job is primarily "don't lose revenue," with growth as real but secondary upside, matching how most CSM roles are actually scoped.

The threshold on the retention piece is what makes this a performance plan rather than a tenure bonus: a CSM whose book drops to 90% GRR earns zero on that component, so there are real stakes tied to churn, usually the thing a SaaS company most needs a CSM plan to prevent.

And because the pay mix is lighter than the AE's (70/30 instead of 50/50), it correctly reflects that a CSM has less unilateral control over a renewal decision than an AE has over closing a new deal, matching pay mix to influence per the first principle in Section 3.

SDR/BDR

- \$85K OTE
 - roughly 72/28 split (\$61K base, \$24K variable)
 - a 20 qualified-meetings-per-month quota (mid-level range per Outbound Sales Pro's 2026 benchmarks)
- \$100 per qualified meeting, no payout below 80% of quota
 - a tiered accelerator (1.5x rate at 100 to 125% of quota, 2x at 125 to 150%)

MEETINGS/ MONTH	% OF QUOTA	MONTHLY VARIABLE	ANNUALIZED (IF SUSTAINED)	% OF TARGET
14	70%	\$0 (below threshold)	\$0	0%
20	100%	\$2,000	\$24,000	100%
27	135%	\$3,150	\$37,800	158%

The key point is that the metric is qualified meetings, not raw dials or emails sent, so the plan is already one step closer to outcome than pure activity.

The 80% threshold and tiered accelerator do the same job as the AE's: no payout for showing up at two-thirds of quota, and real, escalating upside for a rep who blows past it, rather than the flat "same rate no matter what" structure that makes a lot of SDR plans feel like a formality once someone clears their number.

Note: This example only guards against underperformance, not against gaming the single metric. A rep could book 27 "qualified" meetings that are low quality just to chase the accelerator. Consider adding a small kicker (for example, an extra amount for any meeting that converts to a real opportunity within 30 days) to close that loophole without breaking the rule of three.

Pre-Launch Stress Test

Now that you've got your plan drafted, it's time to test and model it.

We recommend running a stress test across four dimensions:

- Cost
- Behavior
- Fairness
- Operations

Our philosophy: a comp plan is ready when it *pays the right amount at target, doesn't blow up on outliers, and still feels fair when messiness shows up.*

1 Cost curve test

Model payout at 0%, 50%, 80%, 100%, 120%, and 150% attainment, then model a realistic team distribution, such as 20% below plan, 60% around plan, and 20% above plan.

Goal: confirm total comp cost stays inside your sales productivity and margin guardrails.

2 Windfall and edge-case test

Run scenarios for an oversized deal, a multi-year or prepay deal, a split-credit deal, a rep with no closes, a deal clawback or cancellation, and a mid-period start or transfer.

Goal: make sure one unusual transaction doesn't create a payout outcome finance can't defend.

3 Behavior alignment test

Check whether the highest-paid deal is actually the best business deal. Compare payout economics across new logo versus expansion versus renewal, one-year versus multi-year, discounted versus full-price, and upfront cash versus delayed cash.

Goal: avoid paying more for revenue quality you value less.

4 Fairness and quota realism test

Pressure-test the plan against different books of business: a strong territory versus a weak one, a ramping rep versus a fully productive one, and an enterprise lumpy motion versus a steady SMB motion.

Goal: verify the plan is hard but attainable, not accidentally punitive.

5 Operational readiness test

Validate that the CRM fields required for payout are populated and auditable, that payout timing is unambiguous, that exception handling is defined, and that finance, payroll, and sales managers can all reproduce results.

Rule of thumb: if three people can't independently calculate the same payout on sample deals, the plan is not launch-ready.

What to require before sign-off



Three to five historical backtests on real deals



One upside scenario and one downside scenario



A dispute checklist for common exceptions



A rep-facing one-pager with example earnings outcomes



Performance-driving comp plans share the following in common:
fewer metrics, thresholds, meaningful upside, and quotas
calibrated to what a rep can actually produce.

The three worked examples in this report show what that looks like
in dollars and cents for an AE, a CSM, and an SDR. And, the
diagnostic and stress test give you a way to check whether your
own plans measure up, and where to focus first if they don't.

You can (and should!) [run your plans and proposed plans through Atlas](#). And if you need a hand, we'd love to walk you through it.

[Book time with our team today](#)



How many AEs do I need to hit \$46M next year?

COMP PLAN DRAFT

AE Compensation Plan — FY26
Mid-Market segment · 14 reps · Designed by Sam K.

PLAN ELEMENTS

- OTE \$160K
- Pay Mix
- Quota-to-OTE
- Accelerators
- Ramp

Ask Atlas
I know this plan inside and out

3 pending decisions

Is our Mid-Market AE plan calibrated for FY26?

Your **7.5x quota-to-OTE** is above the **5-6x** Mid-Market benchmark, and you're missing accelerators above 100%.

SUGGESTED EDIT
7.5x quota-to-OTE
5.5x quota-to-OTE

Accept