

FOR PE & VC PARTNERS

Sales' *Moneyball* problem.

Every revenue team in your portfolio measures outputs. The ones that scale measure the right *input*. Here is the one most CROs don't.

Nate Bliss

Chief Sales Officer, Kinsta • ~30% ARR CAGR over six years

There is one metric your portfolio CROs are almost certainly not measuring. Not because it's exotic. Because it's boring. And because measuring it honestly exposes what's broken.

I spent the last six-plus years as Chief Sales Officer at Kinsta, compounding revenue at roughly 30% per year in a crowded managed hosting category. That is roughly double the growth rate of the managed hosting market over the same period. We beat the field on one simple principle: we measured the input, not just the output.

This piece is a field note for PE and VC partners. If you sit on boards of B2B SaaS companies in the \$5M to \$75M ARR band, or you're evaluating one, the pattern below is almost certainly showing up somewhere in your portfolio.

The *output* trap

Every SaaS board deck looks the same. Win rate. MRR. Net retention. Bookings vs plan. Forecast accuracy. Pipeline coverage. CAC payback.

These are outputs. They are lagging indicators. They tell you what already happened. By the time win rate drops two quarters in a row, the damage is six months old and the people who caused it have already left.

In baseball terms, these are RBIs and batting average. They describe the past. They do not coach the present.

The breakthrough in Moneyball was not that the A's found smarter outputs. They found the input that predicted the output. Then they coached it relentlessly.

On-base percentage was the insight. Strike zone discipline was the coachable behavior underneath it. Everything else was noise.

The sales *OPS*

The equivalent of on-base percentage in a B2B sales org is **talk time**. Specifically: the total minutes each revenue-carrying seat spends in qualified buyer-facing conversation in a given week.

Not dials. Not emails sent. Not activities logged. Not meetings booked. Actual, qualified, buyer-facing talk time.

At Kinsta, once we measured it honestly across AEs, AMs, and SDRs, the correlation to bookings was strong enough that we treat it as causation. The reps who put in the talk time produced the revenue. The ones who didn't, didn't. The variance was not subtle.

That sounds obvious when you say it out loud. It is almost never measured in the companies I see.

Why the *pattern* repeats

Three reasons revenue teams skip the input and obsess over the output.

- **It's uncomfortable to measure.** Talk time exposes underperformance on day seven, not quarter-end. Managers who protect favorites fight the metric.

- **CRMs aren't built for it.** Activity tracking measures volume (dials, emails). It rarely captures qualified conversation time. Most teams approximate with logged meetings, which undercounts the real signal.

- **Boards don't ask for it.** Pipeline coverage and forecast accuracy dominate the QBR. No one asks the CRO what the average qualified talk time per seat was last week.

What to ask on *your next* board call

You don't need a new dashboard. You need one question that exposes whether the team is coaching the input.

"What was the average qualified buyer-facing talk time per revenue seat last week, and how did it compare to the prior four weeks?"

If the CRO can answer in 10 seconds, they're measuring the right thing. If they have to come back to you, they're not. And if the answer includes "dials" or "activities," they're measuring volume, not conversation.

Who you *hire* matters more than what you measure

Measuring the right input only works if the team around you can operate on it. The hardest mistake most GTM leaders make, in my experience, is hiring mirrors instead of complements.

My right hand at Kinsta is an analytical systems operator. Same work ethic I have, completely different wiring. She operationalizes what I see in pattern. I trust her with the mechanics the way Billy Beane trusted Paul DePodesta. I bring the framework, she builds the machine.

My best recent hire is a quiet generalist. Broad surface area, low ego, extraordinary execution consistency. The market doesn't know how to price a person like that. I did. He is load-bearing in ways that do not show up in any org chart. He is my Chad Bradford, the player the market misreads because his value doesn't live in the obvious stat.

Most scaling SaaS companies hire for the résumé the board recognizes. The ones that compound hire for the gap in the leader's own skillset. Those are different decisions.

The *whiteboard* pitch

If you put me in front of a portfolio company CEO for an hour, here is what I'd whiteboard:

- Define the input metric that predicts revenue in this specific motion. Talk time is the right answer for most inside-sales SaaS. It is not the only answer.

- Instrument it. Measure it weekly. Publish it to the team.

- Audit the current hiring criteria for mirror bias. Identify the one role where a complement would compound the founder's or CRO's capacity.

- Make RevOps the engine. Recruiting is the fuel. The team on the field is the product.

These are not novel ideas. They are disciplined ones. Most portfolio GTM problems I see are not strategy problems. They are measurement and personnel problems dressed up as strategy problems.

Want this *applied* to a portfolio company?

If the pattern above matches what you're seeing in a portfolio company right now, I'm opening a small amount of advisory bandwidth in 2026 for exactly this kind of engagement.

Schedule a Conversation