What Private Equity Keeps Getting Wrong About Talent

TL PARTNERS | JULY 2026 | THOMAS IGEME

Are you happy with your executive failure rate?

Not the industry's rate. Yours. The CROs replaced in year one who didn't fix revenue. The reorgs that moved boxes and not numbers. The CEO who was "not the scaling profile" until his successor hit the same wall, at which point the problem turned out to live somewhere other than the corner office.

If you're happy with that rate, stop reading. This paper has nothing for you. If you're not, I want to show you three places where the standard playbook misreads the problem, because I have sat on both sides of it: I built the portfolio talent function at a roughly $4B software PE firm from a standing start, and I now run a strategy advisory for tech companies and the funds that back them.

Misread one: you cannot tell a people problem from a system problem

Start with what funds do well. You diligence financials to the second decimal. You model the thesis under three downside cases. Then you assess the management team that has to deliver all of it with a few partner dinners and a reference call, and after close, when the numbers wobble, you reach for the two tools you have: replace an executive or reorganize.

Here is the uncomfortable mechanic underneath: those two moves treat different diseases, and almost nothing in the standard toolkit tells you which disease you have. When work drifts from strategy, the cause is either WHO (the wrong executive for this thesis) or HOW (an operating model that would defeat any executive). Replace the CRO when the real constraint is a broken handoff between sales, marketing, and customer success, and you burn twelve months onboarding a successor into the same broken system. Reorganize when the real constraint is one seat, and you put the whole company through change management to avoid one hard conversation. Either mistake costs a year, and most holds only have five or six of those.

The steelman: assessment firms exist, and the good ones are genuinely good at the WHO question in general. What they produce is a point-in-time read of a person against a generic scorecard. What they do not produce is the WHO question scored against your specific thesis, and they are structurally blind to HOW, because you cannot learn how work flows through an organization by interviewing individuals about themselves.

The fix is not a better scorecard. It is two instruments run on the same frame: an executive assessment that answers WHO, and a strategy alignment audit that answers HOW, both scored against the WHAT that only you can set, the thesis. Run both and the expensive question answers itself before you spend a year on the wrong fix. In my portfolio years, the diligence version of this ran on four acquisitions before close, and the pattern held every time: the teams looked different once you scored them against the actual plan instead of against "good executive."

The obvious objection: what about when both are broken, a mediocre executive inside an operating model that would defeat a great one? That is the most common finding, not the edge case, and it is why the pairing matters. The instruments do not just distinguish, they sequence. You cannot read WHO clearly through a broken HOW, because a strong leader inside a system that routes decisions nowhere looks exactly like a weak one. So the order is: repair the system enough to see the person, then judge the person against the thesis. Funds run it in the opposite order every week, and pay for a search they may not have needed.

Misread two: you are assessing executives for the last economy

Two things changed underneath the standard executive profile, and most diligence hasn't noticed either one.

First, holds got longer. When the average hold stretches toward six, seven, eight years, you can no longer buy your way past a weak bench at exit speed. Development and succession stop being HR hygiene and become the only way the team you close with is still the right team at exit. A fund that cannot build leaders is now structurally dependent on recruiting luck, twice per seat per hold.

Second, and bigger: the composition of work itself is changing. Work is now produced by people, by AI, and by combinations of the two that shift quarter to quarter. That moves the executive bar. The question is no longer only "has she scaled a revenue org," it is "can this team redesign who does what, including the parts that stop being human." I have led AI enablement inside enterprises: the leaders who carry it are not the most technical, they are the ones who can hold a team to one set of priorities while the ground moves.

So here is the stake: within one fund cycle, leadership diligence that does not read a management team's capacity to restructure work around AI will be malpractice. Not because AI is fashionable, but because a 2026-vintage thesis already assumes margin structure that only exists if the team can do that redesign. You are underwriting it whether you assess it or not.

Misread three: your comp chart tells the talent partner what they are

Look at how your fund pays its operating talent seat, and you will find the honest org chart.

If talent is truly a value-creation lever, the seat should be built like one: in the deal room for diligence, accountable for a thesis-critical system, and carrying economics that look like the value it creates. The firm where I ran talent matched operating-side carry to deal-side carry. That single design choice did more than any mandate: it told me, the partners, and every portfolio CEO exactly what the function was for. Value creation, not support.

Most funds do the opposite and then wonder why the function acts like HR. Pay a seat like a cost center and it will behave like one: compliance, process, and searches outsourced to whoever the deal partner already knows. You get exactly the talent function you paid for.

The test is one question at your next partner meeting: does our head of portfolio talent see deals before the check is written, and do they have carry? If the answer is no twice, you do not have a talent strategy. You have a staffing function with a strategy-shaped title.

What alignment actually buys you

Underneath all three misreads sits one idea, and it is the flag this firm plants: alignment is not consensus, and it is not a static state you achieve at the January offsite. Alignment is a rate. It is how fast an organization reconciles the work people actually do to a WHAT that keeps changing, because in any real hold, the WHAT changes: a pivot, an add-on, a new platform bet, an AI reset.

Measured that way, alignment is the leading indicator that buys you the asset every fund actually wants and none can underwrite today: adaptive capacity. A perfectly aligned organization that reconciles slowly is one strategy change away from mediocre. An organization that reconciles fast compounds through every change you throw at it. The gap between the two shows up in behaviors months before it shows up in the numbers, which means it is measurable while it is still cheap to fix. We measure it. That is the practice.

One question to leave with, the same one this paper opened with, now with its edge showing: your executive failure rate is not bad luck, and it is not a talent shortage. It is the predictable output of a toolkit that cannot tell WHO from HOW and has not noticed that both questions changed. Fix the toolkit and the rate moves.

We present this argument, with the instruments behind it, at fund partner meetings. Fifteen minutes of material, your questions after.