The reporting tax

Every leadership team pays to find out how the work is going. Most pay twice, once in the reporting and once in the surprise.

POINT OF VIEW 7 PAGES · 15 MINUTE READ FOR CEOS, CHROS, OPERATING PARTNERS, PORTFOLIO TALENT LEADERS
AUGUST 2026 THOMAS IGEME

THE ARGUMENT IN ONE PARAGRAPH

Reporting exists because the numbers arrive too late to steer by. It is the price of seeing earlier, and at most companies that price has risen quietly until the seeing costs more than the steering is worth. The fashionable response is to take the human out of the loop and read the artifacts. That trades one failure for another. Artifacts record what happened and hold no record of what anyone intended, so there is nothing to compare them against. An instrument that predicts has to carry both terms at weekly cadence, and it has to make the second one nearly free.

01

OBSERVED

Two options. Both of them fail.

The control system you cannot afford

Every board pack, every weekly business review, every status roll-up exists for one reason. The numbers that finally tell you whether the plan is working arrive after the quarter they describe. Nobody wants to steer a company from the rearview mirror, so leadership installs reporting to see earlier. That reporting is a tax, it is paid in the time of the people doing the work, and at most companies nobody has repriced it in years.

I watched this from inside a software private equity firm, where I built the portfolio talent function from a standing start. The operating thesis was faster laps: get portfolio companies iterating and reporting on a weekly rhythm so the fund could see problems while they were still cheap. The instrument was a weekly performance spreadsheet. Getting it filled in every week took an enormous amount of energy, and the complaint that came back from operators was blunt. Half their job had become filling in reports about the work rather than doing it.

The reporting ate the lap

That is the shape of the trap. We asked for faster laps, and the reporting ate the lap. The thing installed to measure the work consumed the time to do it. A control system that degrades what it measures has become an overhead line that produces a document.

The cost is real money and it is easy to size. Take a 500-person company at $100,000 fully loaded per head. That is $12.5M of payroll a quarter. If internal reporting has crept to even a tenth of the time, and at most companies I have seen it is higher among exactly the senior people whose judgment you are paying for, that is **$1.25M a quarter spent producing a description of what already happened.** Same arithmetic as the drift cost in The Portfolio Talent Operating System, applied to the instrument this time.

Directional. Order of magnitude from headcount and loaded cost. Run it with your own numbers; the point survives a wide range.

NOT ALL OF THE BURDEN WAS WASTE

The sheet had a second effect, and I did not see it for years. Once a week it made a person stop and answer why this number matters. That question is most of alignment compressed into a sentence. Entropy pulls an organization off its stated direction continuously, and something has to pull back at a comparable rate. The sheet did that. It did it expensively and it did it badly, and it did it.

So the design question is which half of the burden to remove. **Being made to think was the part that was working. Being made to transcribe a week you already lived was the part that was not.** Almost every reporting system in existence asks for a great deal of the second and almost none of the first, which is how a ritual ends up both exhausting and uninformative.

So here is the strategic problem, stated at the altitude a CEO or a general partner actually holds it: **you cannot govern on lagging indicators, and you cannot afford leading ones.** Either break that trap or price quarterly surprise as a permanent cost of ownership.

02

OBSERVED

The tax, and the question of who pays it.

Why reporting decays

Reporting rituals decay for a structural reason, and laziness has nothing to do with it. The person filling in the weekly sheet spends their Friday afternoon reconstructing a week they already lived, and the output travels upward to someone else. They get back no decision, no help, and no time. The tax has a payer and it has a beneficiary, and they are different people. Every reporting system with that property decays, and the decay is rational.

The second stage is worse. As the burden rises the quality falls, because people stop writing to inform and start writing to discharge an obligation, which produces green status on dying work. By the time a reporting ritual is universally complied with and universally resented, it has become a machine for manufacturing false confidence.

THE GOALS-PLATFORM PATTERN

The modern version of this is the goals platform. A company buys one, cascades objectives into it, and then discovers that the weekly or monthly or quarterly updating of those goals never saturates. I have heard this diagnosed inside a large consumer-software company as historical cultural nonchalance about goal setting, which is a fair description of the symptom and a poor description of the cause. The cause is that updating a goal is a tax with a decayed collection rate.

The important thing about that diagnosis is what it rules out. The root cause is not that people lack visibility into what matters. In most of the rooms I work in, people can tell you the top three priorities without looking them up. What they will not do, week after week, is pay to maintain someone else's picture of their work.

THE CONCLUSION MOST LEADERS DRAW HERE IS WRONG

Faced with a decayed reporting ritual, the usual read is that the organization is not ready for instrumentation, and that discipline has to come first. That gets the order backwards. Low update discipline is the symptom the instrument treats. The test is whether the instrument reduces the reporting burden or adds to it, and almost every product in this category adds to it.

03

EMERGING PATTERN

One detailed account plus our own field reads.

Why listening does not replace asking

If the tax is the problem, the obvious move is to stop taxing. Remove the human from the loop and read the exhaust: the commits, the tickets, the calendar, the message traffic. Nobody fills anything in and the signal flows for free. That is the promise behind every passive observation and sentiment product in the market, and the appeal is obvious.

It runs into two problems, and the first one is empirical. In a recent conversation with the senior people leader of a large consumer-software company, I heard the cleanest statement of it I have come across. His company is remarkably quiet on public channels. Most real conversation happens in direct messages, which no tool is going to read. And where public channels do get used, they get used for two things: broadcasting information, and arguing with the direction. He gave an example of a recent policy decision that drew a fifty-message thread, overwhelmingly hostile, from the part of the company most inclined to object. Point a sentiment engine at that and you conclude the company is in revolt, when what you have actually measured is the self-selection of who posts.

Empty, or unrepresentative

So the channel is either quiet or skewed, and from inside the tool you cannot tell which. A silent channel and a healthy one produce the same data, and no amount of tuning separates them.

The second problem is structural and it is the deeper one. Passive observation holds one term. It can see what happened. It has no record of what anybody said they were going to do, so it has nothing to compare the outcome against. An artifact stream tells you a team shipped four things last week. It cannot tell you those were the wrong four, because it never knew which four were promised.

04

MECHANISM

Laid out so you can interrogate any piece.

Both terms, every week

What predicts is the distance between the plan and the artifacts, measured often enough to have a trend. A team that says it will do five things and does three has told you something. A team that does that for six weeks running has told you a great deal more, and none of it was available in either term alone.

That is a calibration record, and calibration is about as good a predictor of future delivery as organizations have. It is also, for what it is worth, how every other engineering discipline decides whether an instrument can be trusted. A quarterly goal platform holds both terms in principle and samples them four times a year, which is a destination and not a lap. A passive tool samples continuously and holds one term. Neither can compute the number.

The loop that produces it is simple to describe. At the start of the week the people doing the work name what matters. At the end of the week the instrument reads how it appears to have gone, drafts a status, and asks that same person to confirm or correct it. That is the whole mechanism, and the design choice that matters is in the last four words.

WHERE THE INFORMATION ACTUALLY LIVES

The obvious objection is that a machine-drafted status confirmed by a human is passive observation with a rubber stamp on it. The answer is that the information was never in the draft. **It is in the correction.** When somebody edits "looks on track" into "stalled, legal has not come back," you have learned the cause, and no artifact anywhere in the stack contains the cause. You have also learned that a named person is now holding it.

This is why the burden question resolves. Reporting asks a person to reconstruct the week. Correction asks them to judge it. The first is authorship and the second is a decision, and the person making the decision is the one who benefits from getting it right, because their reality reaches the top without them writing anything.

The fund cannot close the gap itself

There is a version of this that stops one step short, and I ran it. We collected what moved and how far. The explaining happened later, upward, in a meeting, and by then it was my account of their week. What I would build now asks for two more things at the point of correction: why the work went where it went, and what the team is doing differently next week. Both are one line. Both belong to the person closest to the problem, and neither is available to anybody else at any price.

That matters more at a fund than inside a single company, and the reason is arithmetic. A portfolio talent leader covering a dozen companies cannot go deep in any of them. Any design that routes the diagnosis to that person fails at the second company, and then quietly fails at the rest. **The version that scales is the one where the people closest to the work narrow their own drift and report that they did it.** The fund's job is to make the gap visible and to ask the question every week. It is not to close the gap, and at portfolio scale it cannot be.

THE FAILURE MODE ON THE OTHER SIDE

Once you can see drift weekly, the reflex is to instrument the cause. A number slips, you find the leading indicator that would have warned you earlier, and you add it. Then another. I did this at the fund and it does not converge, because each new metric describes one path the work took away from the plan, and there are always more paths. That is whack-a-mole with a reporting cost attached, and it re-raises the tax the loop was installed to cut. **Adding a metric is what a fund does instead of asking somebody why.**

It also gives you two things nobody currently has, and they are worth separating because one is much better evidenced than the other.

The first is a validity check. If nothing in a team is ever corrected, either the machine read is excellent or the loop is dead, and you can tell those apart by whether corrections cluster in some teams and vanish in others. Any roll-up sitting on top of an uncorrected loop should be treated as unverified. That is worth more than it sounds: no existing reporting system can tell you whether its own inputs are real. **Measure correction rate. Completion rate is the vanity metric that produced the dead goals platform in the first place.**

The second is a prediction, and it is the part we are still testing. Every correction moves a draft in a direction. Someone either revises it up, because the work went better than the artifacts suggested, or down, because it did not. A team that revises downward week after week is telling you that its visible output is running ahead of its real progress, which is the shape a miss makes while it is still forming. Our working claim is that a run of downward corrections shows up weeks before the number does. We have design-partner data pointing that way and not enough of it to be sure.

WHAT IS EVIDENCED HERE, AND HOW WELL

STRONG SIGNAL

Reporting rituals decay

Directly observed and repeatedly evidenced, inside a fund and across advisory clients. The payer and the beneficiary are different people, and the ritual erodes on schedule.

EMERGING PATTERN

The channel is quiet or skewed

Multiple indicators that are not yet triangulated: one detailed practitioner account, plus our own reads in client Slack estates. Worth testing in your own before trusting any sentiment product.

HYPOTHESIS TO VALIDATE

Downward corrections lead a miss

Plausible and explicitly unproven. Design-partner data only, small sample, short window. The validity use of correction rate rests on firmer ground than the predictive one, and this is the claim in this paper I would most like to be wrong about in public.

05

APPLICATION

One free test, then four seats.

What to do on Monday

Before any of this becomes a purchase, there is a test that costs an afternoon and tells you whether the argument above even applies to you.

THE FREE VERSION: COUNT THE TERMS YOU HOLD

Pick the three priorities that matter most this quarter. For each one, ask two questions. Is there a written record of what the responsible team said it would do this week? And is there a record of what actually happened? Most companies find they hold the second in some rough form and the first only quarterly, if at all.

If you hold one term, or neither, no product will give you a drift reading, because there is nothing to subtract. That is the finding. It also tells you which half you are missing, and in my experience the missing half is almost always the first one, which is the cheap one to fix.

Two terms are worth defining here, because the four seats below use them and neither is standard vocabulary. **Realignment half-life** is the number of weeks between a strategy change and the work visibly moving to match it. **A blind spot** is a priority showing green with nobody reporting real work against it. Both come out of The Chief Work Officer, and both are computable only once you hold both terms.

What follows is split honestly. The first move in each seat needs no tool and no budget. The second assumes you have decided to instrument.

Chief executive

This week. You are the only person who can cut the tax, because retiring a report is an executive act and nobody below you will do it. Before you approve any new instrument, name the report it replaces. One in, one out, enforced at your staff meeting. If nothing can be named, the answer is that you do not want the instrument.

Once instrumented. Ask for realignment half-life and a blind-spot count, weekly, by priority. The first time you hear a priority is dead should not be the quarterly review.

Chief people officer

This week. Stop reporting goal completion to the executive team and report the share of goals that changed in the last thirty days. You can pull that from the platform you already own. A goal nobody has touched in a quarter is not a goal anybody is using, and that one chart will tell your executive team more about the health of the ritual than any completion percentage ever has.

Once instrumented. Make the loop serve the manager first, so adoption is pulled by usefulness and your function never becomes the nag. Report correction rate.

Operating partner

This week. Add one question to diligence and one to your next board meeting. In diligence: what share of the senior team's week produces internal reporting? At the board: who assembled this pack, and how many hours did it take. Almost nobody asks either, both are answerable on the spot, and at the arithmetic in section one it is a value-creation lever hiding in plain sight.

Across the portfolio. Mandate the question and leave the method alone. Every company answers half-life and blind-spot count however it chooses to get there. Mandating a tool spends political capital you will want later.

Portfolio talent leader

This week. List your companies and, from memory, write down how long it took you to hear about the last real surprise at each one. Weeks, or months. That ranking is a coverage map, it costs you twenty minutes, and it is usually uncomfortable in a useful way.

Once instrumented. Use the per-company read as triage for your own capacity, which is your actual constraint. Read it per company. Companies that share no single goal cannot share a score, and a blended one hides the company that needs you.

The discipline that protects your capacity. Ask for the why and the change alongside the number, and resist writing either one yourself. The moment the diagnosis becomes your work product, your coverage ceiling becomes the portfolio's ceiling.

06

BOUNDARY

Where this earns nothing.

When this does not apply

Three conditions rule this out. An instrument with named edges is easier to trust than one sold as universal.

There is no stated strategy. The whole mechanism compares work against a declared intent. Where leadership cannot state one, no instrument helps, and the honest first move is to say so.

Leadership can still personally read the work. In a company small enough that the leaders sit inside every thread, the room is the instrument, and it is a better one. Past the first management layer that stops being true.

Nothing important is broken. Some companies are genuinely in a state where the things that matter are pushes to continue and not problems to fix. An early-warning system earns little in that state, and it is worth being honest that the value concentrates around inflections: new ownership, a reset, a new chief executive, an integration, a platform bet. The moment briefs in the library begin with those inflections.

> **AND ONE CONDITION THAT LOOKS DISQUALIFYING AND IS NOT**
>
> A collapsed goal process is the most common reason leaders rule this out, and it is the weakest one. The failure mode to avoid is stacking a second reporting surface on top of a first that nobody maintains. If the loop replaces the ritual, the collapsed process is the problem being solved. If it joins the ritual, you have raised the tax and you should not proceed.

If you want to know what your reporting actually costs and whether the work has drifted from the plan, the two-week Strategy Alignment Audit is the read that tells you.

Fixed fee, about 20 minutes from each of one or two senior leaders, every finding labeled by confidence. We build portfolio talent operating systems for PE firms, and a lot of the work happens inside specific portfolio companies, doing the org and talent build directly. The continuous version of this loop is **Alyn**, which runs in Slack and is in design-partner pilots now. Reporting rituals and passive observation both fail as ways of seeing the work. The argument they fail against is **Return on Work**; the method that replaces them is **The Weekly Loop**; who owns the comparison is **The Chief Work Officer**. How the read itself runs: **Anatomy of an Audit**. Full library at **tlpartners.biz/library**.

thomas@tlpartners.biz · **tlpartners.biz**