


What Work Costs Now

Every company is learning to watch what AI costs. Almost none is watching what the spend is for. When agents produce the work, the bill becomes the first honest trace of where effort actually pools.

POINT OF VIEW · 4 PAGES · 9 MINUTE READ · FOR CEOs, CFOS, OPERATING PARTNERS
FIRST ARGUED JULY 2026 · THIS VERSION AUGUST 2026 · THOMAS IGEME

THE RETURN ON WORK THESIS



Six papers, one argument. The hub is [Return on Work](#). Full library at tlpartners.biz/library.

THE ARGUMENT IN ONE PARAGRAPH

When agents produce a growing share of the work, the cost of producing work becomes variable, metered, and instant. The market has noticed half of that. Cost tooling is racing to cover AI spend, and every tool in it compares spend to budget. What nobody owns is the comparison between the spend and the strategy: whether the metered dollars are pooling on the bets leadership actually made. That comparison is how [return on work](#) gets priced. This paper asks what the work now costs, and whether the money is following the bets.

01

The denominator started moving

OBSERVED

The assumption every management instrument inherited.

For as long as anyone has managed anything, the cost of work was the slowest number in the building. Payroll moved once a year. Hiring took quarters. Capacity was headcount, and headcount was a planning cycle. Every instrument of management inherited that assumption without noticing it. A budget is a headcount plan. A productivity ratio is output over a number that barely moves. An operating review is a conversation about the numerator, because the denominator was never the part that moved.

The denominator moves by the hour

That assumption just broke, and it broke on the same day the work itself did. When agents produce a growing share of the work, the cost of producing work becomes variable, metered, and instant. A team lead with a corporate card can triple a team's capacity to produce output in an afternoon, and triple the spend again before finance closes the month. The whiplash is already on the record. Uber burned through its entire 2026 AI tools budget by April, and ServiceNow exhausted its annual token allocation before mid-year. **The denominator now moves faster than the annual instruments watching it.** For the first time in the history of management, what work costs is as dynamic as what work produces, and nothing in the standard toolkit was built for a denominator that moves by the hour.

The market has noticed half of this. Cost tooling is racing to cover AI spend, and finance teams are standing up dashboards of tokens, seats, and model bills. Every one of those dashboards compares spend to budget, which answers whether you are overspending and stays silent on whether the spend is buying the plan. What nobody owns is the comparison between the spend and the strategy, and that comparison is the subject of this paper.

02

EVIDENCE

The leaderboard case, and what each drift signature now costs.

Misallocation shows up in the bill before the P&L

The hub's core observation was that the gap between strategy and work shows up in behaviors months before it shows up in the numbers. The cost side has a sharper version of the same claim. **Misallocation shows up in the bill before it shows up in the P&L**, because the bill is itemized weekly and the P&L is a quarter behind.

Drift is now a variable cost

In the human-only era, misalignment burned salary-time. That loss was real and invisible. The payroll ran either way, no ledger line moved, and the cost of a team spending six weeks on a priority that had quietly died never appeared anywhere a CFO looks. Agents change the accounting. A misaligned team with agents burns metered dollars doing the wrong thing faster, and the waste has a line item. Drift stops being a cultural cost and becomes a variable cost.

The clearest case is the one that ended the leaderboard era. Meta made AI-driven work results a formal performance requirement in November 2025, and by April 2026 an internal leaderboard ranked eighty-five thousand employees by token consumption. Consumption rose twenty-two percent in a month. No output signal moved with it, and the company pulled the leaderboard within days of it leaking. Rank people by tokens and you get tokens. The mistake was the unit of measure, consumption standing in for work.

WHAT EACH DRIFT SIGNATURE COSTS NOW

Pooled effort The tokens pool exactly where the effort does, including on a priority that quietly died.	Outlived experiment A meter still running on a plan whose decision date has passed.	Agent sprawl Three teams on three subscriptions doing one category of work.	Velocity without direction The most expensive signature of all, and now astonishingly affordable per unit.
--	--	--	---

One implication lands outside the people function entirely. For twenty years the alignment conversation lived with the head of people, filed under culture. A variable cost line moves it to the CFO's staff meeting. When misalignment shows up as spend, the question of whether the work is following the strategy stops being a soft question, and the seat accountable for it stops being optional.

03

MARKET MAP

The two ends are covered. The middle is not.

Who owns the spend question today

The honest survey looks like the one in the hub, one level down. Spend to budget is owned. Spend to strategy is not.

WHERE THE SPEND QUESTION SITS

Declaring intent

NO DOLLARS ATTACHED

The goal platforms cascade priorities. **None of them knows what a priority costs to pursue this week.** Engineering intelligence comes closest, tracing human effort to initiatives, but effort was a proxy for cost in an era when cost meant time, and it does not meter the agent half of the work, which is exactly the half that moves.

Spend to strategy

UNOWNED

Whether the metered dollars are pooling on the bets leadership actually made. **Today this is a person:** someone senior squints at the AI line in the monthly review and asks what we are getting for it, and the room answers with anecdotes.

Metering spend

MATURING FAST

FinOps practice, cloud cost platforms, and the AI vendors' own dashboards report spend by team, by model, by seat, and by hour. **Every one of them compares spend to budget.**

A perfectly efficient spend on the wrong priority passes every cost review ever written: the rate was negotiated, the usage was optimized, the budget was met, and the quarter was wasted.

Accenture ran the whole arc inside a year. Employees were told AI use would factor into advancement, then found the company rationing tokens spent on basic tasks, and in leaked internal audio a strategy lead conceded that spend had become unpredictable while leadership at the CFO, COO, and CIO level was still asking whether the company was getting value for it. The largest consultancy on earth, whose product is answering such questions, is asking this one of itself in a meeting. That is not a tooling gap. It is an unowned seat.

How you measure return on work

The instrument is the weekly loop, carrying a second column. The method is [The Weekly Loop](#), and the cost column attaches to the same cadence: the week's work-production spend, cut by stated priority rather than by cost center. The strategic cut of the bill is by bet, because teams are how the org chart thinks and bets are how the strategy does.

Two numbers fall out, and they pair with the two the loop already produces.

Allocation fidelity

The share of work-production spend landing on stated priorities, read weekly. It is the cost twin of the drift reading, and what matters is the trend and the size of the unexplained remainder.

Realignment cost

The price of the lag the hub measured. Realignment half-life is defined in [Return on Work](#), and realignment cost is what that lag burned.

Allocation fidelity will never hit a hundred in a healthy organization, because real weeks contain fires and maintenance, and a plan that leaves no room for them is fiction. Realignment cost is the spend stranded on the old direction while the organization turned. Between the two numbers, a leadership team can finally say what a slow pivot costs in dollars, and what an instrumented one saves.

Nothing about this adds a meeting. The read rides the weekly loop the team already runs, and what reaches leadership is one page with four numbers: drift per priority and realignment half-life from the loop, allocation fidelity and realignment cost from this one. And it needs an owner. The Chief Work Officer holds the comparison. Finance supplies the

ACCENTURE, LEAKED
INTERNAL AUDIO

04

METHOD

Two numbers riding the
weekly loop a team already
runs.

Cut the bill by bet

denominator and keeps its own discipline. The people function does not lose the alignment file to the CFO. It finally gets the file funded.

The blind spot gets a cost face too, and it is the most expensive object in this whole argument: the priority the board believes is on track with no real work against it, sitting next to a fat spend line pooling where nothing was declared. Undeclared exploration with a meter running. The goal tool cannot see it, the cost tool cannot see it, and the two of them side by side cannot see it either, because seeing it requires the comparison neither one owns.

BOUNDARY CONDITIONS, NAMED OUT LOUD

One guardrail belongs in the design, stated plainly. The unit of account is the bet. Read spend and drift at the level of priorities and teams, and keep the instrument away from individual surveillance. The moment it is pointed at people it stops reading the organization and starts deforming it. The confirmation step that keeps the work signal honest is the same step that keeps the cost signal humane.

05

FORWARD CLAIM

Three tradeoffs that become decidable with a number per bet.

The decisions it unlocks

The point of measuring return on work is the tradeoffs it makes decidable. A COO I trade notes with put it flatly: the game is becoming ROI, and the real skill is trading off among the efficient ways to get it. That skill needs a number per bet, and today nobody has one. With one, three decisions change character.

Routing

Which work goes to people, which to agents, and which to blends stops being ideology in either direction and becomes a return comparison per priority, revisited as capabilities shift.

Exploration

Declared exploration gets a budget and a decision date, and the meter makes the date real. The distinction between declared and undeclared exploration becomes enforceable the moment exploration has a spend line.

Concentration

When the return per bet is visible, the oldest portfolio question, are we spread across more bets than we can fund attention for, finally gets asked with evidence at the level of a single company's quarter.

The AI bill is now a strategy artifact. Where the spend pools is a confession about where the effort pools, which makes it the cheapest honest read of a company there has ever been. The fund-side version of that read, per company and in diligence, is [Reading a Portfolio](#).

The seat that owns this comparison may take a decade to appear on an org chart. The bill will not wait even that long. It is already itemized, the spend is already elastic, and the only open question is whether anyone reads it against the plan. **Companies that do will know their return on work in week two. Everyone else will get the news from the P&L, months late, at metered prices.**

This is the pricing spoke of the thesis: what the work costs once agents produce part of it, and whether the money is following the bets.

The argument this paper follows from is [Return on Work](#). Read next: [The Weekly Loop](#), which is the cadence this cost column attaches to, and [Reading a Portfolio](#) for the fund-side view. Full library at tllpartners.biz/library.

thomas@tllpartners.biz · tllpartners.biz