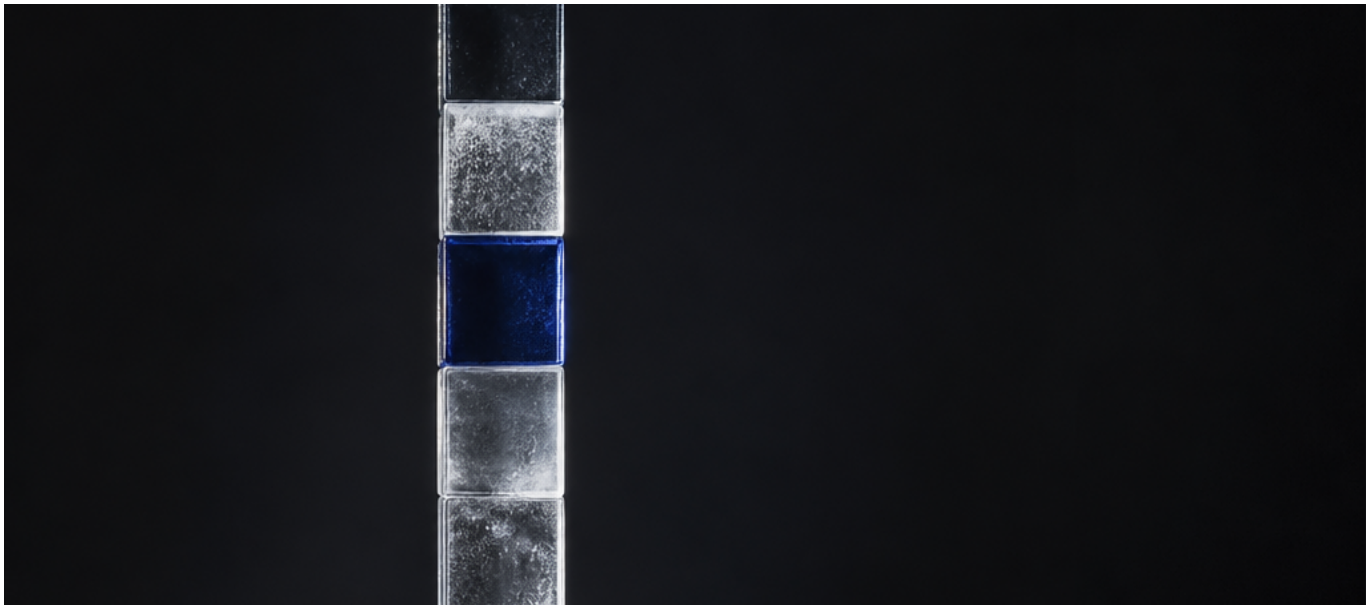


# The 580x Gap: What the State of AI Report 2026 Actually Says About Enterprise AI Spend

In August 2026 the median company spent $12.50 per employee per month on AI. The top 1% spent $7,205. Same vendors, same list prices, 580 times apart.

THE SHORT VERSION

The ninth State of AI Report, published by Nathan Benaich and Air Street Capital on October 8, 2026, puts the top 1% of US companies at a median $7,205 per employee per month on AI in August 2026, against $12.50 for the median firm. That is roughly 580 times, measured across 21,559 firms by Ramp. Meanwhile 69% of the S&P 500 disclose live AI deployments and 2% disclose metrics tracked over time.

The report runs 244 slides and most of the coverage went to model capability. The slide that should worry anyone running a budget is the one built from Ramp card data, where AI

spend was joined to headcount records for 21,559 US firms. It is the rare adoption dataset that does not depend on executives self-reporting their own maturity.

## Why the adoption gap is not about price

A seat is a seat. A 580 times difference in per-head spend means the top 1% are running something structurally different: many more seats, far more inference, far more of the expensive long-horizon work. The median firm bought a few licenses and stopped.

What makes the Ramp cut useful is the headcount link. The top third of adopters, spending $33.67 per worker per month, added 10.2% headcount over two years, including 12% growth in entry-level hiring. Light adopters showed no separation from the control group at all. Not worse, just indistinguishable.

Two caveats the report itself flags, and I would weight them heavily. Most of the employment gains concentrated in technology companies, so this may be a tech-sector story wearing an economy-wide costume. And spend correlates with growth without card data establishing the direction of causation. Fast-growing companies have budget. Budget buys AI.

There is a second limit worth stating plainly. Ramp sees card transactions, not outcomes. A company spending $7,205 per employee per month could be running a disciplined agent fleet or burning money on shelfware at scale. The dataset separates spenders from non-spenders, and it does not separate good spenders from bad ones. Nothing in the report closes that loop.

## Sixty-nine percent say yes, two percent can prove it

This is the number I would put in front of a board. 69% of S&P 500 companies disclosed live AI deployments. 2% disclosed metrics tracked over time. The other 67% have announced something they are not measuring, or are measuring privately and have chosen not to say so, which for an outside reader amounts to the same epistemic position.

I have written before about [how badly the standard enterprise AI ROI numbers get quoted](https://www.arturmarkus.com/enterprise-ai-roi-5-numbers-everyone-quotes-wrong-from-95-fail-to-mckinseys-37/), and this is the cleanest illustration of why. The

failure-rate debate is mostly an argument between people who did not instrument anything. When 98% of the largest public companies in the US cannot show a tracked series, any claim about AI's aggregate return, bullish or bearish, is an inference about an unmeasured population.

My take: the 2% figure is a better leading indicator than the 69%. Deployment disclosure is cheap talk for an earnings call. A tracked metric over time means someone owns a baseline, which means someone can be wrong in public. I expect that gap to narrow slowly, because closing it requires admitting what the baseline was.

## OpenAI eats its own output at 97.9%, individual users at 0.7%

The Codex usage numbers make the same shape appear inside a single product. In the preceding 28 days, 97.9% of OpenAI staff used Codex. Among organizational ChatGPT users it was 17.3%. Among individual users, 0.7%.

| CODEX USAGE | OPENAI STAFF | ORG CHATGPT USERS | INDIVIDUAL USERS |
|---|---|---|---|
| **Used in past 28 days** | **97.9%** | 17.3% | 0.7% |
| **Share of engineer output tokens (June 2026)** | n/a | **26.8%** | n/a |
| **Share of legal-user output tokens (June 2026)** | n/a | 1.9% | n/a |
| **Assigning 8-hour-scale tasks, Dec 2025** | n/a | n/a | 2.1% |
| **Assigning 8-hour-scale tasks, May 2026** | n/a | n/a | **25.6%** |

That last pair is the most interesting movement in the dataset. The share of individual Codex users handing the model at least one task on an eight-hour scale went from 2.1% in December 2025 to 25.6% in May 2026. In five months, delegating a full working day of work went from a fringe behaviour to a quarter of users. That is a change in what people believe the tool is for, and it happened faster than any procurement cycle I know of.

Inside organizations the composition is shifting too. Since February 2026, weekly active enterprise Codex users grew 108 times in legal, 41 times in sales and recruiting, 26 times in marketing, against 5 times in engineering. Those multiples come off near-zero bases, so

they flatter badly. The honest reading sits in the token share: in June 2026, Codex produced 26.8% of the average organizational engineer's output tokens and 1.9% of the average legal user's. Legal is growing fast and doing almost nothing. Engineering is growing slowly and doing a quarter of the work.

MY READ

*Headcount-based adoption metrics are now actively misleading. "We rolled out to legal" and "legal produces 1.9% of its output through the tool" are both true, and only one of them is a result. If I were setting an internal KPI for 2027, it would be token share or task share by function rather than seat count, precisely because seats are the metric that makes a 108 times chart possible on a base of almost nothing.*

## $519B of funding is chasing an unmeasured outcome

On the supply side the numbers are not ambiguous. AI companies raised $519B through mid-September 2026, per the Dealroom and Air Street dataset. US firms took $393B of it, 76% of the total. China took $59B. Europe took $40B, which is less than a third of the single largest round in the set: OpenAI's $122B.

Concentration is extreme. The 15 largest rounds totaled $312B, and generative-AI companies accounted for $261B of that. So roughly 60% of all AI venture funding this year sits in fifteen cheques. I covered the quarterly shape of this in the [Q3 2026 funding breakdown](https://www.arturmarkus.com/q3-2026-ai-funding-by-the-numbers-103-3b-85-in-154-mega-rounds-and-europes-75-record/); the annual view makes the same point with less noise.

Revenue is following, at least at the top. Combined OpenAI and Anthropic annualized run rate reached roughly $105B by late summer 2026, up from about $30B at the start of the year. That is a 3.5 times move inside eight months, and it is the strongest argument against the pure-bubble reading. Someone is paying, a lot, and increasingly fast.

The uncomfortable pairing is this: $519B of capital and a $105B run rate sit on top of a buyer population where the median firm spends $12.50 per head per month and 98% of the S&P 500 publishes no tracked outcome. The revenue is concentrated in the same place the spend is concentrated, in the top percentile. Whether the median ever converts is the open

question, and nothing in 244 slides answers it.

## Cost per capability is falling 13 times a year

One number cuts in the other direction and deserves equal weight. Epoch AI's estimate in the report puts the cost of reaching a fixed benchmark score falling 47% per quarter since 2023, across five benchmarks. Compounded, that is roughly 13 times per year.

So the median firm's $12.50 buys dramatically more capability than $12.50 bought in 2024. A laggard waiting two more years will acquire the same capability for about 1/169th of today's cost. That is a real argument for patience, and I think it is the strongest one available to a CFO who does not want to move.

My take: it is also wrong for most businesses, and the Ramp headcount data is why. What the top third bought was two years of organizational rewiring, workflow redesign, and the measurement discipline that lets you tell a working deployment from an expensive one. Model access is getting cheap on schedule. That other part does not deflate 47% per quarter. Compute gets cheaper; knowing where to point it does not. A company arriving in 2028 with cheap tokens and no instrumentation will look exactly like the light adopters in this dataset, statistically indistinguishable from the control group.

The specific detail I keep returning to is the entry-level hiring figure. Top-third adopters grew entry-level hiring 12% over two years. The dominant narrative says AI eliminates junior work first. In this dataset, the companies spending most on AI hired more juniors. I do not know whether that holds outside the technology sector, where most of the employment gains concentrated, and the report does not claim it does. It is the single finding most likely to be quoted backwards over the next year, so read the slide rather than the summary.

Three things I would do with this data if I were setting next year's AI budget. Instrument first and publish internally, because joining the 2% costs nothing but the willingness to see a flat line. Measure output share by function rather than seats, since Codex at 26.8% of engineer tokens and 1.9% of legal tokens is the difference between a deployment and a logo on a slide. And stop treating $12.50 per head as prudence. It is the spend level the data says produces no separation at all.



WRITTEN BY

Artur Markus

Software engineer in Basel building AI infrastructure and agentic systems, with a background in pharmaceutical automation, GxP compliance, and validation. I help teams turn AI from a chatbot into a reliable operational layer.

Book a meeting →