


AI • TECH

# Microsoft reports are exposing AI's real cost problem: Using the tech is more expensive than paying human employees

By Jake Angelo  
News Fellow

Add us on

May 22, 2026, 12:56 PM ET



Microsoft CEO Satya Nadella  
SVEN HOPPE/PICTURE ALLIANCE VIA GETTY IMAGES

Firms today are pushing employees to use as much AI as possible to squeeze out the technology's productivity gains. But that pressure is leading to cracks, and those cracks may be irreparable.

Microsoft has reportedly begun canceling most of its direct Claude Code licenses, according to *The Verge*, instead moving engineers toward using GitHub Copilot CLI. That comes just six months after the firm first opened up access to Claude Code, encouraging thousands of its developers, project managers, designers, and other employees to experiment with coding. The tech became popular fast. Perhaps too popular. The scale at which employees use it is now prompting the firm to reverse course on a tool its own engineers had come to rely on. Canceling Claude Code licenses won't affect Microsoft's Foundry deal, which includes investing up to \$5 billion in Anthropic and giving Foundry customers access to Claude models, as well as Anthropic's \$30 billion commitment to purchase Azure compute capacity, according to *The Verge*.

Microsoft isn't the only company scaling back its internal AI use. Uber's CTO Praveen Neppalli Naga told *The Information* in April that the firm had already burnt through its entire 2026 AI coding tools budget in just four months. That comes after the company had actively incentivized adoption through internal leaderboards ranking teams by AI tool usage.

The reports may throw cold water on the bets tech's biggest firms have placed on the technology. While some cling to the promise of an AI "renaissance" or "revolution," the cost of adoption is proving a stubborn bottleneck. These developments also suggest that the economics of replacing or augmenting human labor with AI may be more complicated than some early forecasts originally implied. That echoes what Bryan Catanzaro, vice president of applied deep learning at Nvidia, recently said in an interview with Axios.

"For my team, the cost of compute is far beyond the costs of the employees," he said.

Anthropic didn't immediately respond to *Fortune's* request for comment. Microsoft didn't provide a comment.

## An emerging AI paradox: cheaper tokens, bigger bills

Uber and Microsoft aren't the only firms pushing employees to use as much AI as possible. Like at Uber, a Meta employee crafted a leaderboard, fittingly named "Claudeconomics," after Anthropic's AI model, to track which workers are using the most AI. Amazon is pushing its employees to "toxiemaxx," or use as many AI tokens as possible (the basic building blocks of AI compute).

But with a token-based pricing system, the work gets more expensive with more use and better efficiency. Goldman Sachs recently forecasted that agentic AI could drive a 24-fold increase in token consumption by 2030 as consumers and enterprises adopt AI agents, reaching a staggering 120 quadrillion tokens per month. As businesses turn to AI agents to boost productivity, aggregate costs could rise sharply even if the price of each token falls.

But as consumption increases, the cost of individual AI tokens is expected to fall sharply. A recent report from research firm Gartner found that by 2030, inference on a one-trillion-parameter LLM—in simple terms, a highly sophisticated AI model—will cost AI firms nearly 90% less than it did in 2025. Even so, Gartner predicted that cheaper tokens won't translate to cheaper enterprise AI because agentic models require far more tokens per task than standard models, increased consumption can outpace falling unit costs, and AI providers won't fully pass through lower costs to consumers. In turn, inference costs are likely to push higher.

"Chief Product Officers (CPOs) should not confuse the deflation of commodity tokens with the democratization of frontier reasoning," Gartner senior director analyst Will Sommer warned in a statement.

That reality may complicate the grand plans some firms have for deploying AI agents. Nvidia CEO Jensen Huang recently said he thinks 100 AI agents will one day work alongside every employee at his company.

Huang is part of a broader wave of CEOs touting an agentic future in which digital workers operate across the enterprise. But if token consumption rises faster than unit costs fall, that future could come with a much heavier bill than executives expect.

In 2001, Fortune first convened the smartest people we know, bringing together CEOs and founders, builders and investors, thinkers and doers. Since then, **Fortune Brainstorm Tech** has been the place where bold ideas collide. From **June 8-10**, we will return to **Aspen**—where it all began—to mark 25 years of Brainstorm. [Register now.](#)


## About the Author

By Jake Angelo  
News Fellow  
[See full bio >](#)

## Latest in AI

4 hours ago

5 hours ago



ASIA • CHINA

**Huawei touts chip breakthrough to shorten gap with TSMC**

By Bloomberg  
May 21, 2026

1 day ago

ECONOMY • DISRUPTION

**'Nobody knows anything': 'this time is different': t phrases that define — and haunt — the AI economy**

By Nick Lichtenberg  
May 23, 2026

1 day ago

**Unlock the reporting leaders trust to stay ahead**

1 day ago

1 day ago

AI • PRODUCTIVITY

**BoFA says you'll be 10x more productive with AI. Ignore the 0.1% result so far**

By Nick Lichtenberg  
May 24, 2026

COMMENTARY • MEDIA

**I was one of the world's first influencers. AI just killed the whole category — and created something better**

By David S. Bennahum  
May 24, 2026

AI • RESEARCH

**AI hallucinations are infiltrating expert work—and entering the permanent body of knowledge**

By Tristan Bove  
May 24, 2026

## Most Popular

4 days ago

9 hours ago

1 day ago

SUCCESS

**Jeff Bezos wants the bottom half of earners to pay zero income tax—he says nurses making just \$75K should save \$12K a year**

By Preston Fore  
May 21, 2026

1 day ago

ECONOMY

**The pig in the python: Baby Boomers are strangling the economy they built by refusing to move or retire**

By Nick Lichtenberg  
May 25, 2026

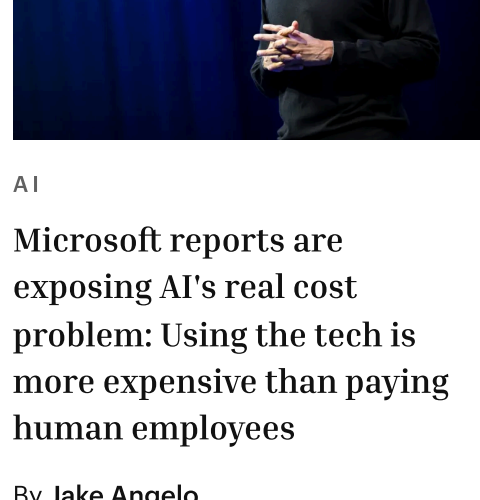
2 days ago

LIFESTYLE

**'freed up' his son from having to get a driver's license—and he's one of many Gen Zers who aren't willing to drive**

By Sasha Rogelberg  
May 24, 2026

3 days ago



REAL ESTATE

**Inside the 'stealth wealth' playbook: How Silicon Valley's elite buy multimillion-dollar mansions without leaving a paper trail**

By Sydney Lake  
May 24, 2026

SUCCESS

**This 39-year-old quit his lineman job during the pandemic and built a \$50 million company in his backyard**

By Nick Lichtenberg  
May 23, 2026

AI

**Microsoft reports are exposing AI's real cost problem: Using the tech is more expensive than paying human employees**

By Jake Angelo  
May 24, 2026

Rankings Sections Customer Support Commercial Services About Us Press Center Work At Fortune Terms And Conditions Site Map

## Cookies On This Site

We and our 487 partners store and access personal data, like browsing data or unique identifiers, on your device. Selecting I Accept enables tracking technologies to support the purposes shown under we and our partners process data to provide. Selecting Reject All or withdrawing your consent will disable them. If trackers are disabled, some content and ads you see may not be as relevant to you. You can resurface this menu to change your choices or withdraw consent at any time by clicking the Manage Preferences link on the bottom of the webpage [or the floating icon on the bottom-left of the webpage, if applicable]. Your choices will have effect within our Website. For more details, refer to our Privacy Policy. We and our partners process data to provide:

Use precise geolocation data. Actively scan device characteristics for identification. Store and/or access information on a device. Personalised advertising and content, advertising and content measurement, audience research and services development. [List of Partners \(vendors\)](#)

I Accept

Reject All

Show Purposes