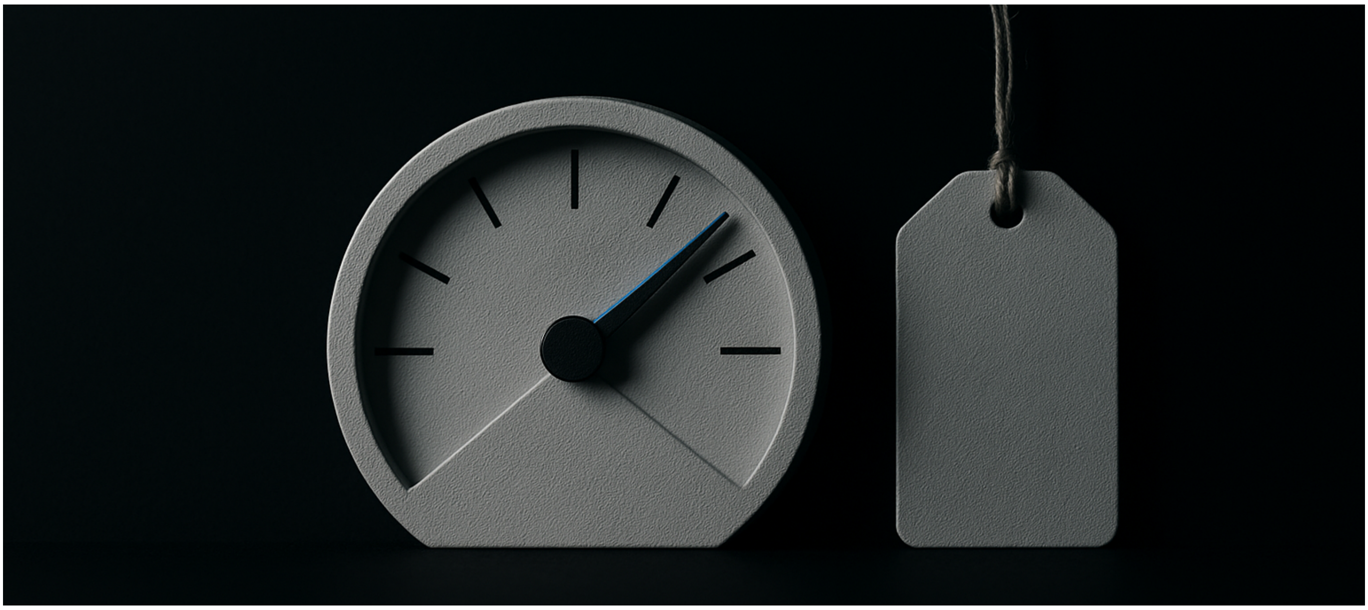

Copilot vs Cursor vs Claude Code: The Honest Cost Comparison After the September 1 Pricing Reset

Nobody raised a seat price on September 1. Copilot Business still costs $19. Your bill went up anyway, because the meter moved instead of the sticker.

TL;DR

On September 1, 2026, GitHub cut included Copilot AI credits by 36.7% on Business (3,000 → 1,900 per user/month) and 44.3% on Enterprise (7,000 → 3,900), while Anthropic's Claude Sonnet 5 introductory API rate expired and moved 50% higher to $3/$15 per million tokens. Seat prices didn't budge. For most teams doing normal feature work, Copilot Business at $19/seat plus measurable overage is still the cheapest floor, and it's my pick — but the reset makes Copilot's cost far less predictable than its price page suggests, and if your workflow is genuinely agentic, Cursor's flat $40 seat is the better bet on variance alone.

The decision: this is the week the math changed

If you signed a Copilot contract in June, you did the math against 3,000 credits per user per month. That was a transitional allowance, and it expired. Per GitHub's usage-based billing documentation, the standard monthly allowance is 1,900 AI credits per user for Copilot Business and 3,900 for Copilot Enterprise. The June 1–September 1 transitional numbers were 3,000 and 7,000.

That is a 1,100-credit drop on Business (-36.7%) and a 3,100-credit drop on Enterprise (-44.3%), as documented in the September 1 credit cliff analysis. One AI credit equals $0.01. So a team that had settled into burning 3,000 credits per user per month now pays 1,100 overage credits × $0.01 = **$11 extra per seat per month** — a 58% increase on the effective cost of a $19 plan, with no change to the advertised price.

The same day, Anthropic's introductory rate for Claude Sonnet 5 lapsed. EnterpriseDNA reported the expiry date as August 31, 2026, moving pricing from $2/M input and $10/M output to $3/M and $15/M — a clean 50% increase on both legs. A workload of 10M input plus 10M output tokens per month goes from $120 to $180.

-44.3%
Copilot Enterprise included AI credits, 7,000 → 3,900 per user/month
GitHub Docs / ampm-aiops.com, September 2026
+50%
Claude Sonnet 5 API rate on both input and output legs
enterprisedna.co, August 9, 2026
$0
change to Copilot Business, Copilot Enterprise, or Cursor Business seat prices
bytetech247.com, August 19, 2026

Two of the three major vendors repriced on the same day, in the same direction, using the same mechanism: leave the sticker alone, move the meter. If your AI tooling budget was set in Q2, your Q4 forecast is now wrong.

Two vendors repriced on the same day without touching a single seat price. The invoice is now the only honest price page.

The contenders

GitHub Copilot

Copilot Business is $19 per seat per month and Enterprise is $39, and those prices did not change with the credit reduction. The structural detail that matters: code completions and next-edit suggestions remain free and unlimited under usage-based billing, and the metered cost applies to Copilot Chat and agent workflows, per Tech Times' reporting on the billing change. Credits are consumed on tokens processed — input plus output plus cached — at each model's published API rate. So Copilot's cost curve is flat until you start using agents, then it tracks the underlying model market directly.

Cursor

Cursor Business is $40 per user per month, roughly $480 per seat per year, including 500 fast premium requests plus unlimited slow requests, according to ModelsLab's pricing comparison. The unit is a request — not a token, not a credit. That matters more than it sounds: a request is a coarse unit, so Cursor absorbs the volatility of token counts inside a fixed envelope. It is more than twice Copilot Business's seat price and it did not move on September 1.

Claude Code

Claude Code access runs through Claude subscription tiers — $20/month for Pro and $100 or $200/month for Max — with heavy enterprise use estimated at $50–150 per developer per month per the same ModelsLab analysis. Underneath, you're exposed to raw token pricing: Sonnet 5 at $3/$15 per million as of September 1, and Claude Opus 5 at approximately $5/M input and $25/M output with a 1M-token context window, launched late July 2026. Claude Code is the only one of the three where the September 1 increase hits as a direct, unbuffered pass-through.

Head to head

Note what's missing: there is no cost-per-task row, because the three platforms meter fundamentally different units — Copilot bills AI credits, Cursor bills fast premium requests, Claude Code bills raw tokens — and no standardized cost-per-task comparison across them was published in August or September 2026. Anyone selling you one built it themselves.

Attribute	Copilot Business	Cursor Business	Claude Code
Seat price / user / month	$19	$40	$20 (Pro) – $200 (Max)
Billing unit	AI credits ($0.01 each)	Fast premium requests	Raw tokens
Included allowance	1,900 credits/user/mo	500 fast requests + unlimited slow	Tier-dependent
Sept 1, 2026 change	-36.7% included credits	None	+50% Sonnet 5 API rate
Free/unlimited tier	Completions + next-edit suggestions	Unlimited slow requests	None documented
Cost exposure to model price moves	Direct (credits track API rates)	Buffered by request unit	Direct pass-through
10-dev annual floor	$2,280 ($190/mo)	$4,800 ($480/seat/yr)	$6,000–18,000 (est., heavy use)
Largest context window offered	Model-dependent	Model-dependent	1M tokens (Opus 5)

!

Watch out The 10-developer annual floor row compares a seat floor against a usage estimate. Copilot Business at $190/month or $2,280/year assumes zero overage; the Claude Code $6,000–18,000/year range is derived from the $50–150 per developer per month estimate for heavy enterprise use. These are not like-for-like, and treating them as such is exactly the error that gets budgets approved and then blown.

The Copilot credit math, done properly

Measure your actual credit burn per developer, not your seat count. Because Copilot credits are consumed on tokens processed at each model’s published API rate, your Copilot bill is a derivative of the model market — and the model market moved on September 1 too.

If your Copilot agent workflows are routed through Sonnet 5, the 50% API increase and the 36.7% credit cut compound: fewer credits, and each credit buys fewer tokens of the same model. Neither vendor announced a price rise. I wrote about this pass-through structure when GitHub first shifted Copilot to token-based billing in June; September 1 is the first time the mechanism has been used to raise effective prices without a price announcement.

The claim

“Copilot Business is $19 per seat, cheapest of the three, easy call for a cost-conscious team.”

The reality

$19 is the floor, not the price. A team burning 3,000 credits per user per month — the level GitHub’s own transitional allowance normalized between June and September — now pays $11 per seat in overage, taking effective cost to $30. Tech Times reported agentic bills jumping 10x to 50x for power users under this model.

That 10x-to-50x figure is the one to sit with. It is not a claim that everyone’s bill multiplied; it’s a claim about power users, and every engineering org has them. Your two or three most agent-native developers can generate overage that dwarfs the seat cost of the other twenty.

Model choice is now a budget lever inside Copilot. GPT-5.6 Sol is listed at promotional pricing of $4.00/M input, $20.00/M output and $0.40/M cached input — a promotion stated as running through at least November 21, 2026, down from legacy $5/$30 rates, per OpenAI’s API pricing documentation and the AWS Bedrock reduced-pricing announcement. Sonnet 5 at $3/$15 is cheaper on both legs than Sol’s promotional rate. Opus 5 at roughly $5/$25 is the most expensive of the three, and it’s the one your developers will reach for on hard problems.

My take

Read “at least November 21, 2026” as a floor, not a commitment. If your Copilot credit budget currently balances on Sol’s $4/$20 promotional rate, you are one announcement away from a repeat of September 1 — same mechanism, different vendor. I would model Sol at its legacy $5/$30 rate for any forecast extending into Q1 2027. That’s my read, not a vendor statement.

Where each one wins

Pick Copilot Business if completions are most of your value

The free-and-unlimited completions tier is Copilot's real moat. If your developers get most of their lift from inline completion and next-edit suggestions, with Chat and agents occasional rather than constant, Copilot Business at $19 plus modest overage is structurally cheaper, because the highest-frequency interaction is unmetered. A 10-developer team at $190/month, $2,280/year, is a floor no competitor here approaches.

Pick Cursor Business if you need budget certainty above all

Cursor is the only one of the three where nothing changed on September 1. The request-based unit — 500 fast premium requests plus unlimited slow requests at $40 per seat — is a coarser meter than credits or tokens, and coarse meters absorb volatility. You pay roughly double Copilot's seat price for a number you can put in a spreadsheet twelve months out.

What works about Cursor's model

- $40/seat did not move on September 1 while two competitors repriced.
- Unlimited slow requests give a genuine unmetered fallback, not just a throttle.
- Request-based billing decouples your invoice from per-million-token API moves.

Rough edges

- 2.1x Copilot Business's seat price — $4,800/year for 10 devs vs $2,280.
- 500 fast requests is a hard number with no published cost-per-task equivalence to credits or tokens, so you cannot pre-compute whether it's enough.
- No documented free unlimited completions tier equivalent to Copilot's.

Pick Claude Code if context window is the binding constraint

Opus 5's 1M-token context window is the only capability in this comparison that isn't a pricing artifact. If your work involves reasoning over large codebases in a single pass — migrations, cross-cutting refactors, legacy archaeology — that window is worth paying for, and roughly $5/$25 per million is the price of admission. The estimated $50–150 per developer per month for heavy enterprise use is $6,000–18,000/year for ten developers, and you should only sign up for it if you can name the specific tasks that justify it.

i

The unit problem You cannot benchmark these three against each other on cost per task using published data. Credits, requests, and tokens are not convertible, and no standardized cross-platform comparison was published in August or September 2026. Any vendor deck showing you one is showing you a model with assumptions buried in it. Run a two-week pilot on your own repos instead.

My pick

For a 10 to 50 person engineering org doing feature work on an established codebase, **Copilot Business remains the right default — but instrument it this month rather than next quarter.**

The reasoning: the unmetered completions tier covers the majority of daily AI interaction for most teams, so the $19 seat does most of the work. The overage exposure is real but measurable — $0.01 per credit is the clearest unit price in this market. And at $2,280/year for ten developers, the floor is low enough that even a meaningful overage bill lands under Cursor's $4,800 floor.

My take

Two failure modes I'd guard against. First, the power-user tail: Tech Times documented 10x-to-50x agentic bill increases for heavy users, and if three of your twenty developers live in agent mode, they will dominate your invoice. Set per-user credit alerts before month-end, not after. Second, model routing drift: because credits are consumed at each model's published API rate, letting developers default to Opus 5 at ~$5/$25 instead of Sonnet 5 at $3/$15 quietly raises cost per task with no config change and no announcement. Both are policy problems, not vendor problems — which is why I'd still choose Copilot and then govern it.

Where I'd switch: if a two-week measurement shows your team consistently above

1,900 credits per user per month with the overage curve still climbing, Cursor's $40 flat seat becomes the cheaper and calmer choice. The crossover point is roughly 2,100 credits per user per month on paper ($19 + $21 overage = $40), but I'd move earlier, because variance costs you more than the delta once you're defending a budget you can't forecast.

And if your problems are genuinely large-context — you can name three recurring tasks that need a million tokens in one pass — run Claude Code alongside Copilot for those specific developers rather than replacing anything. Anthropic has been repricing in both directions this quarter; I covered how cache-read price cuts broke working prompt patterns, and I expect the same volatility that raised Sonnet 5 by 50% will move something else next quarter. Keep the Claude exposure small enough that either direction is survivable.

What to do in the next two weeks

Pull actual credit burn per user, not per team

Copilot Business now includes 1,900 credits/user/month and Enterprise 3,900. Find out who is above and by how much. The distribution matters more than the average, because overage is charged at a flat $0.01 per credit.

Reprice your Q4 forecast at the post-September-1 rates

If your model assumed 3,000 Business or 7,000 Enterprise credits, it is wrong by 1,100 and 3,100 credits per user per month respectively. Add $11/seat/month for any team that had normalized to the 3,000-credit allowance.

Recompute any direct-API workloads at $3/$15

Sonnet 5 pipelines built against the $2/$10 introductory rate cost 50% more as of September 1. A 10M-in/10M-out monthly workload moved from $120 to $180. Check anything with a token budget baked into a business case.

Set a model-routing default and enforce it

Sonnet 5 at $3/$15, Sol at $4/$20 promotional, Opus 5 at roughly $5/$25. Pick your default, make the expensive model an explicit opt-in for named task types, and revisit before the Sol promotion's stated November 21, 2026 window closes.

Bottom line

Seat prices are now marketing; allowances and API rates are the price. Copilot Business at $19 is still the right default for most teams because completions stay unmetered — but the 36.7% credit cut and the 50% Sonnet 5 increase landed on the same day, and if you haven't measured per-developer credit burn since August, your budget is already wrong. If you want a second pair of eyes on the numbers before you renew or switch, book a call →

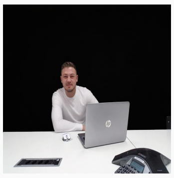
WRITTEN BY
Artur Markus
Software engineer in Basel building AI infrastructure and agentic systems, with a background in pharmaceutical automation, GxP compliance, and validation. I help teams turn AI from a chatbot into a reliable operational layer.

Book a meeting →