

Your \$30 Copilot Seat Is Now a Cover Charge

Microsoft 365 Copilot against going direct to OpenAI or Anthropic: list prices, admin controls, Microsoft 365 reach, and what 500 seats cost at three levels of use

David Soden

Independent analyst and practitioner, enterprise automation and applied AI
davidsoden.com/market-assessment

About this report

Every substantive claim in this report carries an evidence classification and, where applicable, a numbered source. Facts are separated from vendor claims, third-party estimates, the author’s assessments, and untested hypotheses. Forward-looking sections use scenarios with observable tripwires rather than point forecasts. Stated assumptions are listed before the analysis begins.

PUBLISHED FOR PUBLIC READING | SHARE FREELY, UNMODIFIED, WITH ATTRIBUTION

© 2026 David Soden, davidsoden.com/market-assessment. All rights reserved. You may share this file complete and unmodified, and quote short passages with attribution. Republishing under another name, excerpting into a commercial deliverable, resale, and use as machine learning training data require written consent. Full terms appear under Rights and Permitted Use on the final page.

Market Assessment reports are published and commissioned at davidsoden.com/market-assessment. This document is analysis, not investment, legal, or procurement advice.

Scope and evidentiary standard

This report tests one argument against the public record: that Microsoft 365 Copilot is the most expensive way for a business to buy AI, that a company could go direct to OpenAI or Anthropic, connect it to Microsoft 365 and pay less for equal or better admin controls, and that Microsoft failed strategically by not building its own model for Microsoft 365. The trigger is Microsoft’s redesigned Copilot app, announced on 25 September 2026, which moves its agent features onto usage-based billing on top of the existing seat.

It covers the Microsoft 365 Copilot enterprise add-on and the Copilot Business add-on for firms of up to 300 users, GitHub Copilot as context, ChatGPT Business and Enterprise, and Claude Team and Enterprise. Google Gemini is out of scope. All prices are US list prices in US dollars as published on the vendors’ own pages on 26 September 2026, before any enterprise agreement discount. Microsoft renamed Microsoft 365 Copilot to Microsoft Copilot in its documentation during 2026; this report keeps the older name because it is the one on most invoices.

The published numbers in this category need discounting for three reasons. Vendor pricing pages are mostly undated and changed several times in 2026 alone. Model names change faster than contracts: Microsoft, OpenAI and Anthropic each shipped new model families between June and September 2026. And the most quoted adoption figures come from paywalled reporting or vendor-commissioned studies. Every substantive claim below therefore carries one of six classifications, and every claim resting on an external source carries a numbered reference that links to the source list. Calculations are the author’s own and are labelled as assessments. Where a number could not be found, the report says so rather than estimating it.

Label	Definition	How to treat it
FACT	Verifiable in the cited primary source: a filing, press release, official	Reliable as stated. Check the source if the exact wording matters.

Label	Definition	How to treat it
	documentation, or standards body announcement.	
VENDOR CLAIM	Reported by a vendor about its own business or product. Self-measured and not audited by any third party.	The vendor said it. Directionally useful, not a measurement. Definitions of the metric are usually undisclosed.
THIRD-PARTY ESTIMATE	A research firm forecast or survey result. Methodology is proprietary and generally not reproducible.	Use for direction and order of magnitude only. Do not build a model on it.
ASSESSMENT	The author's reasoned judgment from the cited evidence. Falsifiable, and the reasoning is shown so it can be argued with.	This is analysis, not reporting. Disagree with it where your evidence differs.
HYPOTHESIS	A proposition offered for testing. Not currently supported by sufficient evidence to assert.	Treat as a research question, not a conclusion. Each one names what would confirm or refute it.
SCENARIO	A structured future state with a stated subjective probability. The probability is the author's judgment, not a calculated figure.	Use the leading indicators attached to each, not the probability. The indicators are observable; the probability is not.

Assumptions this analysis rests on

Five premises are assumed rather than shown. Each names the conclusion that fails if it is wrong.

One. List prices stand in for what buyers pay

Enterprise agreements, Azure consumption commitments and negotiated OpenAI and Anthropic contracts all discount list. The cost comparison uses list on every side. If one vendor discounts much further than the others for a given buyer, the ranking in the cost section changes for that buyer.

Two. The Copilot seat covers everyday chat without throttling

Microsoft says GPT-5.6 and Sonnet 5 are included in the seat and Opus 5 is included with limits. The cost model treats medium chat use as covered by the seat. If Microsoft throttles medium users or moves more chat onto credits, the finding that Microsoft is cheaper at medium use fails.

Three. Microsoft's own claim about Cowork cost is roughly right

Microsoft publishes no credit rate per Cowork task. The model prices Microsoft's agent work using Microsoft's own claim that Cowork is 30 to 40 percent cheaper than Anthropic's Claude Cowork. If that claim is false, Microsoft falls to rough parity with Claude Enterprise at heavy use rather than beating it. The report shows both cases.

Four. The token volumes describe real work

The three usage profiles are constructed, not observed. Agent tasks in particular vary by an order of magnitude in token use. A firm whose agent tasks are ten times larger than assumed here should rerun the numbers, and the direction of the heavy-use result becomes more sensitive to the unpublished Microsoft credit rates.

Five. The firm keeps its Microsoft 365 suite either way

Going direct to OpenAI or Anthropic does not remove the Microsoft 365 E3, E5 or Business licence underneath. The comparison is therefore the incremental cost of AI on top of a suite the firm already pays for. A firm leaving Microsoft 365 altogether faces a different question that this report does not answer.

The call

ASSESSMENT Half right. For a firm whose staff mostly chat or leave the tool unused, Microsoft is the dearest way to buy AI, because its seat costs more than a direct seat. At medium and heavy use the order flips: the seat already covers chat that a direct buyer pays for by the token, and Microsoft claims its agent work costs less than Anthropic's own. The strategic charge fails. Microsoft pays OpenAI nothing per answer, no longer pays it a revenue share, and is already building its own models behind one credit meter. It was late, not wrong. What would overturn the cost half: Microsoft publishing credit rates for Code and Autopilot that price agent work above the direct API.

My judgment on the evidence assembled here, nothing more. There is data I can't see and data I may have missed. New evidence can move me, including to the opposite position, and I reserve the right to move.

Executive summary

The September 25 announcement did not replace the \$30 seat. It turned it into an entry fee. Chat stays inside the seat, while Cowork, Code, Autopilot and the newest frontier models are metered in Copilot Credits on top. That changes the buying question from whether the seat is worth \$30 to what a given user will cost once the meter runs, and Microsoft has not yet published enough to answer it for two of the three new features.

FINDING 1 The post that prompted this got two facts wrong.

FACT The \$30 Microsoft 365 Copilot seat is still listed and Microsoft says nothing in it changes. GitHub Copilot is not folded into the new app: Microsoft's own post says developers keep using GitHub Copilot, and GitHub is changing its seat billing, not ending it. What converged is the meter, not the products: GitHub AI Credits and Copilot Credits both cost \$0.01. [\[1\]](#) [\[2\]](#) [\[3\]](#) [\[13\]](#) [\[16\]](#)

FINDING 2 The seat is now a cover charge, and the drinks have no price list yet.

FACT Microsoft publishes \$0.01 per Copilot Credit and a full rate card for Copilot Studio agents, but no credit amount per Cowork task, per Code app or per Autopilot hour. From 2

November 2026, new Copilot Business licences bought through resellers arrive with pay-as-you-go switched on. ^{[5] [7] [9] [10] [11]}

FINDING 3 Microsoft is the most expensive option only for light and idle users.

ASSESSMENT On the author's model at 500 seats, Microsoft costs \$30.00 per user per month against \$22.52 for Claude Enterprise at light use, but \$37.02 against \$42.14 at medium use and \$82.65 against \$112.34 at heavy use. At a plausible mix of users the two land within \$110 a year of each other. ChatGPT Enterprise cannot be totalled because OpenAI does not publish its seat price. ^{[3] [5] [32] [34] [35]}

FINDING 4 Direct vendors match Microsoft on identity and spend controls and trail it on compliance.

ASSESSMENT SSO, SCIM, no-training defaults and hard spending caps are documented on all three. Native DLP, retention and eDiscovery are not: Microsoft Purview covers ChatGPT Enterprise for audit and eDiscovery but not for DLP or sensitivity labels, and Anthropic relies on third-party DLP partners and offers no EU data residency. ^{[26] [37] [41] [49]}

FINDING 5 Both direct vendors now reach every Microsoft 365 surface in the brief.

FACT Anthropic's Microsoft 365 connector reads and, where an admin allows, writes across SharePoint, OneDrive, Outlook and Teams, and its Excel, PowerPoint and Word add-ins are generally available. OpenAI has live apps for SharePoint, Outlook and Teams and one add-in for Word, Excel and PowerPoint, but no Outlook add-in. ^{[28] [29] [30] [31] [42] [43]}

FINDING 6 Microsoft pays OpenAI nothing per token.

FACT Since April 2026 Microsoft no longer pays OpenAI a revenue share, while OpenAI keeps paying Microsoft one through 2030. Microsoft's CEO calls its OpenAI model licence royalty free through 2032. The only reported outside model bill is Anthropic, on pace for about \$500 million a year in January 2026 according to The Information. ^{[54] [58] [61]}

FINDING 7 The own-model critique is about timing, not direction.

ASSESSMENT Microsoft's AI chief says the OpenAI contract capped Microsoft's own frontier training until late 2025. Since then Microsoft has shipped seven MAI models, including one tuned for Excel, reorganised Copilot around cost of goods, and moved every new feature onto one credit meter. That is most of the strategy the argument says Microsoft missed. ^{[65] [67] [69] [70]}

FINDING 8 The 7 percent is real and it is rising fast.

FACT Microsoft reported over 30 million paid Copilot seats in July 2026 against over 450 million commercial Microsoft 365 seats, about 6.7 percent. Paid seats doubled between January and July. When workers can choose, a Recon Analytics survey found only 8 percent pick Copilot. ^{[56] [57] [58] [81]}

What Microsoft announced, and what the post got wrong

FACT Microsoft's 25 September post, by Jared Spataro, describes a Copilot app with three sections. Home is "where Chat and Cowork come together in one place." Code is "powered by the

same underlying technology as GitHub Copilot" and runs in a sandbox that can be hosted in the customer's tenant. Autopilot, "previously called Scout," is an always-on agent that "is expanding to private preview at the end of the month." Home and Code start in Microsoft's Frontier early-access programme "in the coming weeks." ^{[1] [77] [78]}

FACT The billing change came in a companion post by Nicole Herskowitz. Usage-based billing covers "Cowork, Code, Autopilot, new advanced experiences in apps like SharePoint, and frontier models like Fable and Astra." These "stay off until an admin creates a spending policy in the Microsoft 365 admin center." On the seat itself she wrote that "there are no changes to the value you already have." Satya Nadella described the shift on the July earnings call as "evolving our business model beyond per-seat to per-seat-plus-consumption." ^{[2] [59]}

Error one: the \$30 seat is not going away

FACT On 26 September 2026 Microsoft's enterprise page lists Microsoft 365 Copilot at \$30.00 per user per month paid yearly, or \$31.50 paid monthly on an annual commitment. Usage billing is added on top of that seat, not in place of it. ^{[2] [3]}

ASSESSMENT The post has the direction backwards. Microsoft did not swap a flat fee for metering. It kept the fee and added the meter. The post's claim that this matches what OpenAI and Anthropic already offer is closer to true than it looks, but for the opposite reason: Anthropic's Enterprise plan is also a seat plus usage at API rates, and OpenAI sells Enterprise on credits or tokens. Seat plus consumption is now the shape of all three offers, and Microsoft moved to it rather than away from seats.

Error two: GitHub Copilot is not folded in

FACT Microsoft's post says "your software developers continue to use GitHub Copilot for their day-to-day work with more connectivity into the Copilot platform." GitHub's own changelog from 28 August says Copilot Business and Enterprise prices are not changing and that upfront per-seat charges start on 1 October 2026, which is a change to how the separate subscription is billed, not an end to it. ^{[1] [16]}

ASSESSMENT What did converge is the unit of account. GitHub moved to GitHub AI Credits at \$0.01 each on 1 June 2026, and Copilot Credits are also \$0.01. The products stay separate. The meter is shared, and Microsoft's reseller notice already lists a GitHub Copilot harness among the services billed through Microsoft 365 pay-as-you-go. ^{[11] [13] [15]}

What the post got right

FACT CNBC reported on 25 September that of more than 450 million commercial Office 365 seats, fewer than 7 percent have Copilot licences. Microsoft's own figures, over 30 million paid seats in July against over 450 million paid commercial seats in January, give 6.7 percent on the author's arithmetic. The post's agent-builder claim is loosely right: Code builds apps from plain language, and Copilot Studio agents join the same cost management in October. ^{[2] [56] [57] [75] [76]}

What is generally available and what is not

Feature	Status on 26 September 2026	How it is billed	Source
Copilot Chat in apps	Generally available	Inside the \$30 seat; GPT-5.6 and Sonnet 5 included, Opus 5 with limits	Herskowitz, 25 Sep
Cowork	Generally available since 16 June 2026	Copilot Credits on top of the seat	Lamanna, 16 Jun
Home (Chat plus Cowork)	Frontier programme rollout in the coming weeks	Chat in seat, Cowork on credits	Spataro, 25 Sep
Code	Frontier first; preview for Microsoft 365 Premium and Pro later this year	Copilot Credits	Spataro, 25 Sep
Autopilot (was Scout)	Private preview from end of September	Copilot Credits	Spataro, 25 Sep
Copilot Studio agents	Generally available	No charge for licensed users within fair use; credits otherwise	Learn, 3 Aug
GitHub Copilot	Generally available, separate subscription	Seat plus GitHub AI Credits	GitHub, 27 Apr

Microsoft's price list, and the part it leaves blank

FACT Microsoft sells two Copilot add-ons. The enterprise add-on is \$30.00 per user per month paid yearly or \$31.50 paid monthly. Copilot Business, for firms of up to 300 users, lists at \$21.00 paid yearly and \$25.20 monthly, with a promotional \$18.00 for annual commitments made between 1 July and 31 December 2026, first year only. Bundled, Business Standard with Copilot is \$23.50 and Business Premium with Copilot is \$32.00. Prices as listed on 26 September 2026. ^{[3] [4]}

FACT Microsoft 365 E7, launched at \$99 per user and generally available from 1 May 2026, bundles E5, Copilot and Agent 365. The suites underneath rose on 1 July 2026: Microsoft 365 E3 from \$36 to \$39 and E5 from \$57 to \$60. Microsoft says the increase does not apply to standalone Copilot SKUs. ^{[17] [18]}

Credits: the rate card that exists

FACT Pay-as-you-go Copilot Credits cost \$0.01 each. Copilot Studio also sells a pack of 25,000 credits a month for \$200, with up to 20 percent off for upfront commitment. Microsoft's Copilot Studio rate card charges 1 credit for a classic answer, 2 for a generative answer, 5 for an agent action and 10 for tenant graph grounding, and marks every row "no charge" for users holding a Microsoft 365 Copilot licence, subject to fair usage limits and excluding computer-using agents. ^{[6] [7] [8]}

Credits: the rate card that does not

FACT For Cowork, Microsoft says the cost of a task is "calculated from four inputs: model use, context retrieval, tool calls, and runtime." It publishes no credit amount for a typical Cowork task, a Code app or an Autopilot hour on the announcement posts, the Cowork launch post or the two Microsoft Learn billing pages. It points buyers to a credit estimator instead. [\[1\]](#) [\[2\]](#) [\[5\]](#) [\[9\]](#) [\[10\]](#)

ASSESSMENT This is the centre of the pricing problem. A finance lead can cap what Copilot spends but cannot forecast it. Microsoft has priced the meter, not the work, and for the two newest features the work is where the money will go.

Who switches the meter on

FACT The first spending policy an admin creates "activates the AI experiences, like Cowork, that are gated by usage-based billing." Policies set tenant, group and per-user monthly limits, and "when users hit the limit, they lose access to agents and services for the rest of the month," although a running task over the per-user limit can finish. An "Auto-apply new services" setting is on by default, so services Microsoft adds later join existing policies automatically. [\[9\]](#) [\[10\]](#)

FACT Two exceptions weaken the off-by-default promise. Microsoft's September partner notice says usage-based billing will be on by default for new Copilot Business licences bought through the reseller channel from 2 November 2026, with a preset monthly limit whose amount the notice does not state. And budgets on the older pay-as-you-go billing for Copilot Chat agents do not stop spending: Microsoft's documentation says the system "doesn't enforce the budget or prevent your organization from exceeding it." [\[11\]](#) [\[12\]](#)

GitHub Copilot tiers, for context

Plan	Price	GitHub AI Credits included per month
Free	\$0	An allowance; 2,000 completions
Pro	\$10 per month	1,500
Pro+	\$39 per month	7,000
Max	\$100 per month	20,000
Business	\$19 per seat per month	1,900 per user, pooled
Enterprise	\$39 per seat per month	3,900 per user, pooled

FACT GitHub replaced premium requests with AI Credits on 1 June 2026 and said plan prices would not change. One credit is \$0.01, code completions stay unlimited on paid plans, and unused credits do not carry over. Figures as listed on 26 September 2026. [\[13\]](#) [\[14\]](#) [\[15\]](#)

What going direct costs

OpenAI

FACT ChatGPT Business Standard costs \$20 per user per month billed annually or \$25 monthly, and a Premium seat with five times the usage is \$100 or \$125. OpenAI positions Business "for teams of 2 to 200 employees" and cut subscription seat prices by \$5 on 2 April 2026. ChatGPT Enterprise is listed as "custom pricing." Prices as listed on 26 September 2026. ^{[19] [20]}

VENDOR CLAIM A figure of about \$60 per user per month for ChatGPT Enterprise circulates widely. The blog most often cited for it traces the number to a single 2023 Reddit post and says OpenAI has never confirmed it. This report does not use it. ^[23]

FACT OpenAI publishes Enterprise token rates. GPT-5.6 Sol, the model Microsoft includes in the Copilot seat, costs \$4.00 per million input tokens, \$0.40 cached and \$20.00 per million output tokens, and Instant usage is billable on these agreements. On the API, gpt-6-astra is \$10.00 input and \$50.00 output, gpt-6-sol \$2.00 and \$10.00, and gpt-5.6-sol's promotional rate runs "at least through November 21, 2026." OpenAI does not publish the dollar price of a ChatGPT credit. ^{[2] [21] [22]}

Anthropic

FACT Claude Team Standard costs \$20 per seat per month billed annually or \$25 monthly, and Premium \$100 or \$125, for teams of 2 to 150. Claude Enterprise is "US\$20/seat/month, billed annually" plus "usage at API rates," with no usage included in the seat. Prices as listed on 26 September 2026. ^{[32] [34]}


FACT Anthropic's help article on Team billing shows an \$80 annual Premium seat, which conflicts with the \$100 on its pricing page. The pricing page is used here. The discrepancy does not affect the 500-seat comparison, which runs on Enterprise. ^[33]

FACT Anthropic's API lists Claude Opus 5.5 at \$4 per million input tokens and \$20 output, Sonnet 5 at \$2 and \$10 after its introductory price became standard, Haiku 4.5 at \$1 and \$5, and Fable 5.1 at \$10 and \$50. Batch processing halves both. Microsoft includes Sonnet 5 in the Copilot seat and bills Fable on credits. ^{[2] [35]}

The size ceiling most comparisons skip

ASSESSMENT The \$20 direct seats quoted in most comparisons do not apply to a 500-person firm. OpenAI positions ChatGPT Business for up to 200 employees and Anthropic caps Claude Team at 150. At 500 seats a buyer lands on Enterprise with both, which means an unpublished seat price from OpenAI and a \$20 seat plus metered usage from Anthropic. In the other direction, a firm of 300 or fewer can buy Copilot Business at \$21, or \$18 on the 2026 promotion, which is the cheapest AI seat of the three vendors.

FIGURE 1 Published AI seat prices, per user per month, annual billing



FACT US dollars, list prices on vendor pages accessed 26 September 2026. Every seat sits on top of a Microsoft 365 suite licence the firm already pays for. ChatGPT Enterprise is omitted because OpenAI does not publish its price. The Claude Enterprise seat carries no usage, so it is not comparable to the others on its own. [\[3\]](#) [\[4\]](#) [\[19\]](#) [\[32\]](#)


Cost at 500 seats, three levels of use

ASSESSMENT The model below is the author's, built from list prices and stated token volumes. It compares Microsoft 365 Copilot with Claude Enterprise because those are the only two options fully priced at 500 seats. Direct usage is priced on Sonnet 5 for chat and Opus 5.5 for agent tasks. An agent task is assumed to read one million input tokens, half of them from cache, and write 30,000 output tokens, which costs \$2.70 at Opus 5.5 list rates. Microsoft's chat is inside the seat. Microsoft's agent tasks are priced at 35 percent below the direct cost, the midpoint of Microsoft's own claim that Cowork is 30 to 40 percent cheaper than Claude Cowork. [\[2\]](#) [\[5\]](#) [\[32\]](#) [\[35\]](#)

Profile	Assumed use per user per month	Direct usage cost	Microsoft per user	Claude Enterprise per user
Light	10 chat turns a working day over 21 days, 4,000 tokens in and 400 out per turn; no agent tasks	\$2.52	\$30.00	\$22.52
Medium	30 chat turns a day, 6,000 in and 600 out; 4 agent tasks	\$22.14	\$37.02	\$42.14

Profile	Assumed use per user per month	Direct usage cost	Microsoft per user	Claude Enterprise per user
Heavy	Medium chat plus 30 agent tasks	\$92.34	\$82.65	\$112.34
Mixed	60 percent light, 30 percent medium, 10 percent heavy	n/a	\$37.37	\$37.39

FIGURE 2 Cost per user per month at 500 seats, author's model



ASSESSMENT US dollars at list prices on 26 September 2026. Constructed, not observed. Microsoft's agent-task cost rests on Microsoft's unaudited claim about Cowork, because Microsoft publishes no per-task credit rate. Code and Autopilot are excluded because their rates are not published. Negotiated discounts and integration costs are not modelled. [\[3\]](#) [\[5\]](#) [\[34\]](#) [\[35\]](#)

ASSESSMENT Across a year at 500 seats, Microsoft costs \$180,000 against Claude Enterprise's \$135,120 at light use, a \$44,880 premium. At medium use it is \$222,120 against \$252,840, and at heavy use \$495,900 against \$674,040. At the mixed profile the totals are \$224,226 and \$224,328. The argument that Microsoft is the most expensive way to buy AI holds only in the first row.

ASSESSMENT The crossover is simple. Microsoft's seat costs \$10 more than Claude's, and it includes chat that Claude bills by the token. Once a user's direct chat bill passes \$10 a month,

which on these token sizes happens somewhere between 26 and 40 chat turns a working day, the Copilot seat is the cheaper one before any agent work is counted.

ASSESSMENT If Microsoft's Cowork claim is wrong and its agent tasks cost exactly what the direct API costs, Microsoft rises to \$40.80 at medium use and \$111.00 at heavy use. That is still below Claude Enterprise in both rows, by a few dollars. Microsoft is not the most expensive option at heavy use under either reading.

THIRD-PARTY ESTIMATE The idle seat is where the argument earns its money. Recon Analytics found that only 35.8 percent of workers with Copilot access use it. Applied to 500 seats, that leaves about 321 idle seats, and at a \$10 monthly gap between an idle Copilot seat and an idle Claude Enterprise seat, about \$38,520 a year spent on nothing. The usage rate is Recon's survey figure, and the dollar figure is the author's arithmetic. ^[81]

ASSESSMENT ChatGPT Enterprise cannot be placed on the chart. At OpenAI's published Enterprise token rates for GPT-5.6 Sol, the same light and medium chat profiles cost \$5.04 and \$22.68 per user per month, double the Sonnet 5 figures, before any seat fee. A buyer negotiating a ChatGPT Enterprise seat anywhere near \$20 would pay more than Claude Enterprise at every level of use on these rates. The model also leaves out three costs that fall only on the direct route: third-party DLP tooling, connector administration, and a second vendor contract. ^[21]

Admin controls, side by side

The table records what each vendor documents on its own pages, not what its marketing claims. Plan names in brackets mark where a control is limited to one tier.

Control	Microsoft 365 Copilot	ChatGPT	Claude
Single sign-on	Entra ID accounts required	SAML SSO (Business and Enterprise)	SSO (Team and Enterprise)
Provisioning	Entra group-based licensing; users already in the directory	SCIM (Enterprise only)	SCIM (Enterprise only)
Audit logs	Purview Audit records each Copilot interaction and model provider; Cowork shows no provider	Compliance Platform logs, 30-day retention (Enterprise)	Audit logs and Compliance API (Enterprise)
Retention	Purview retention policies for Copilot interactions	Admin-set retention (Enterprise); zero retention by contract	Custom retention, 30-day minimum, indefinite by default (Enterprise)
Training on business data	Prompts, responses and Graph data not used to train foundation models	No, by default	No, by default

Control	Microsoft 365 Copilot	ChatGPT	Claude
DLP and compliance	Purview DLP native; label blocking generally available, prompt blocking in preview	Purview covers audit and eDiscovery but not DLP or labels; 18 listed partners	Third-party partners including Forcepoint and Cyberhaven; Purview not named by Anthropic
Spending caps	Credit limits stop usage; older agent budgets alert only	Per-seat credit limits (Business); budgets and overage limits (Enterprise)	Organisation and per-user monthly limits
Roles	AI Administrator and AI Reader Entra roles	RBAC (Enterprise only)	Custom roles (Enterprise only)
Data residency	EU Data Boundary, but Anthropic models excluded	10 storage regions; inference in EEA, US, UAE (Enterprise)	US or global only; no EU option documented

FACT Sources for the Microsoft column are its Learn pages on privacy, audit, retention, DLP, Entra roles and usage billing, dated between June and September 2026. The ChatGPT column draws on OpenAI's pricing page, enterprise privacy page, Compliance Platform and data residency articles, and Microsoft's own page on Purview support for ChatGPT Enterprise. The Claude column draws on Anthropic's pricing page and its help and privacy articles. [\[10\]](#) [\[19\]](#) [\[24\]](#) [\[25\]](#) [\[26\]](#) [\[27\]](#) [\[32\]](#) [\[36\]](#) [\[37\]](#) [\[38\]](#) [\[39\]](#) [\[40\]](#) [\[41\]](#) [\[44\]](#) [\[46\]](#) [\[47\]](#) [\[48\]](#) [\[49\]](#) [\[50\]](#)

ASSESSMENT On identity, training terms and spend caps the direct vendors are equal to Microsoft, and on caps OpenAI and Anthropic are cleaner, because Microsoft still runs one billing path whose budgets only send email. On compliance they are not equal. A Microsoft shop gets DLP, retention and eDiscovery for Copilot inside tools it already runs. Going direct means a Compliance API feed into a separate DLP product, and for ChatGPT, Microsoft's own Purview page says DLP and sensitivity labels are not supported.

How far OpenAI and Anthropic reach into Microsoft 365

Surface	ChatGPT	Claude
Outlook mail	Outlook Email app searches and retrieves mail and drafts replies; sends plain text from shared or delegated mailboxes. No Outlook add-in.	Microsoft 365 connector searches mail; sends and manages drafts if an admin enables writes; cannot attach files. Outlook add-in in beta.
Outlook calendar	Reads events; creates, updates and cancels on supported calendars	Reads meetings; creates, updates and deletes events if enabled
Teams	Sends and replies to messages, reads transcripts; admin sync read-only (Enterprise)	Reads channel and chat messages; sends if enabled. Calendar read-only.

Surface	ChatGPT	Claude
SharePoint and OneDrive	Live search of what the user can open; admin sync index (Enterprise); Purview-encrypted files not indexed; 100 MB cap	Searches and analyses documents; creates and updates files if enabled
Excel, Word, PowerPoint	One add-in for all three; edits Word in place, builds workbooks and slides; no memory or plugins	Separate add-ins for Excel, Word and PowerPoint, all generally available; Word edits with tracked changes

FACT Both vendors' Microsoft 365 access is delegated, requires an Entra administrator to grant tenant consent, and runs through Microsoft Graph. OpenAI notes that a broad Outlook search "may return incomplete results." Anthropic's connector says plainly that "Claude can't attach files to the drafts it creates." ^{[28] [29] [30] [31] [42] [43]}

ASSESSMENT For reading, drafting and editing, the direct route now reaches every surface the brief named, which was not true a year ago. The gap that remains is depth rather than reach. Copilot grounds on the tenant's Graph index by default. The direct tools search live or maintain their own index, and both depend on Graph permissions that Microsoft sets. A buyer is renting access to Microsoft's data plane either way.

What Microsoft pays for the models

FACT The October 2025 restructuring left Microsoft with about 27 percent of OpenAI, extended its rights to OpenAI models and products through 2032, and kept OpenAI's revenue share to Microsoft in place. OpenAI contracted to buy an incremental \$250 billion of Azure services. ^[53]


FACT The April 2026 amendment went further: "Microsoft will no longer pay a revenue share to OpenAI. Revenue share payments from OpenAI to Microsoft continue through 2030." Microsoft's licence to OpenAI IP became non-exclusive. On the April earnings call Nadella described "a frontier model, royalty free, with all the IP rights," and CFO Amy Hood spoke of the IP "as royalty free with the elimination of our rev share to them." ^{[54] [58]}

FACT Microsoft's FY2026 10-K reports \$24.1 billion of revenue from commercial arrangements with OpenAI, including revenue-share payments. It attributes a \$2.6 billion, 12 percent rise in Productivity and Business Processes cost of revenue to "investments in AI infrastructure to support Microsoft 365 Copilot seat and usage growth." It does not disclose a cost or margin for Copilot. ^[55]

THIRD-PARTY ESTIMATE Anthropic is the one outside model bill that has been reported. The Information reported in September 2025 that Microsoft would pay to reach Anthropic's models through AWS, and in January 2026 that Microsoft's Anthropic spending was on pace to hit \$500 million a year. The Information also reported in May 2026 that OpenAI's revenue-share payments to Microsoft are capped at \$38 billion through 2030. None of these figures is confirmed by either company. ^{[60] [61] [62]}

FACT Money also flows back from Anthropic. In November 2025 Anthropic committed to buy \$30 billion of Azure compute and Microsoft committed to invest up to \$5 billion in Anthropic. Microsoft's Q4 FY2026 release records a \$3.2 billion gain on that investment. Since 7 January 2026 Anthropic has operated as a Microsoft subprocessor under Microsoft's data protection terms. ^{[45] [56] [63]}

FIGURE 3 Microsoft Cloud gross margin by quarter, fiscal 2026



FACT As reported by Microsoft's CFO on each earnings call. Microsoft Cloud includes Azure and commercial Microsoft 365, so the line cannot isolate Copilot. It shows direction only. ^{[57] [58] [59]}

ASSESSMENT Microsoft's markup on model access cannot be calculated from anything it discloses. What can be said is where the cost sits. For OpenAI models, the marginal cost is Microsoft's own compute, not a payment to OpenAI. For Anthropic models, it is a real cash cost, reported at a few hundred million dollars a year, partly recycled through Anthropic's Azure commitment. Hood told analysts in July that "model diversification also is a margin improvement opportunity," which is Microsoft saying the model bill is a lever it intends to pull. ^[59]

Phi, MAI, and the model Microsoft did not build

FACT Microsoft Research's Phi models were built to be small. Phi-1, released in June 2023, had 1.3 billion parameters and was trained in four days on eight A100 chips. Phi-3-mini, at 3.8 billion parameters, was pitched in April 2024 as a model that runs "locally on your phone," and Phi-4, at 14 billion, followed in December 2024 under an MIT licence. A 3.3 billion parameter derivative, Phi Silica, runs on Copilot+ PCs and powers on-device rewrite and summarise in Word and Outlook. No Phi-5 has been released. ^{[71] [72] [73] [74]}

FACT Microsoft AI's own models started later. MAI-1-preview, announced in August 2025, was trained on about 15,000 H100s for "certain text use cases within Copilot." In November 2025 Mustafa Suleyman formed an MAI Superintelligence Team. In March 2026 Nadella merged

consumer and commercial Copilot and moved Suleyman to models, writing that Microsoft would build "models that have real product impact, in terms of evals, COGS reduction." [64] [66] [67]

VENDOR CLAIM On 2 June 2026 Microsoft AI launched seven models, including MAI-Thinking-1 and MAI-Code-1-Flash. Microsoft says MAI-Thinking-1 is preferred to Sonnet 4.6 in its own blind evaluations and that its "MAI tuned model for Excel matches GPT 5.4 while being up to 10x more efficient." Neither claim has been independently tested. [65]

Was a Microsoft 365 model realistic?

FACT Suleyman told CNBC in April 2025 that Microsoft's strategy was "off-frontier," to "play a very tight second, given the capital-intensiveness of these models." In November 2025 he told Fortune that the compute cap in the old OpenAI contract was "a big limitation" for a company of Microsoft's scale, and in June 2026 he told VentureBeat that Microsoft was "only sort of set free from our contract with OpenAI about six months ago." The October 2025 terms still apply compute thresholds if Microsoft uses OpenAI IP to develop AGI. [53] [68] [69] [70]

ASSESSMENT A Microsoft 365 model below the cap was always allowed, and Phi shows Microsoft could build small models well. The reason it did not bet Copilot on one is economic. Until April 2026 Microsoft had royalty-free access to OpenAI's frontier models, so a home-built model had to beat free. That calculation changed only when Anthropic's models became a cash cost and OpenAI's licence went non-exclusive, and Microsoft's model push dates from exactly then. Epoch AI's estimate that frontier training costs have grown 2.4 times a year since 2016 explains why Microsoft chose to trail rather than lead. [54] [84]

HYPOTHESIS The brief proposes that a Microsoft 365-specific model with one Copilot and one bill would have locked rivals out of the enterprise. It would be confirmed by evidence that OpenAI and Anthropic won Microsoft 365 customers mainly on model quality that a tuned Microsoft model could have matched. It would be refuted by evidence that they win on product preference and access to Graph, which Microsoft publishes to all comers. The Recon Analytics preference data and the connector documentation above both point toward refutation, but neither tests it directly. [42] [81]

Where the argument is wrong

Microsoft does not pay OpenAI a margin on each answer

FACT The claim that Microsoft could have "kept the margin it now pays out to model vendors" misreads the contracts. Microsoft pays OpenAI no revenue share and no per-token fee it has disclosed, and describes its OpenAI licence as royalty free. The revenue share runs from OpenAI to Microsoft. The only outside model payment on record is Anthropic, reported at an annual pace of about \$500 million, which is small next to the \$24.1 billion Microsoft booked from OpenAI in fiscal 2026. [54] [55] [58] [61]

Microsoft is building the model, late

ASSESSMENT The claim that Microsoft "didn't" build its own model was true until 2025 and is no longer true. Microsoft now has MAI models in production, a CEO memo tying them to cost of goods, a CFO calling model choice a margin lever, a stated model system that separates the agent layer from any one model family, and a single credit meter across Copilot, Code, Autopilot, Copilot Studio and GitHub. What Microsoft chose not to do is make its own model the only one. Copilot runs OpenAI and Anthropic side by side, with more labs promised. [\[2\]](#) [\[59\]](#) [\[65\]](#) [\[67\]](#)

Most expensive only for light and idle seats

ASSESSMENT On list prices at 500 seats, Microsoft is the most expensive option for light users and the cheapest of the two fully priced options for medium and heavy users. For firms of 300 or fewer, Copilot Business is the cheapest AI seat on the market. A blanket claim that Microsoft is the most expensive way to get AI does not survive the arithmetic.

Equal admin controls holds for identity and spend, not compliance

ASSESSMENT Going direct does get equal SSO, provisioning, training terms and spend caps. It does not get equal DLP, retention or eDiscovery without buying and wiring a separate compliance product, and for EU buyers Anthropic documents no EU residency at all. The "equal or better" claim is right for an IT director and wrong for a compliance officer.

The claim that Copilot is a worse wrapper has no test behind it

ASSESSMENT No public, independent benchmark compares Copilot's agent layer with ChatGPT or Claude on the same work. The strongest evidence is preference: Recon Analytics found 70 percent of workers with a free choice pick ChatGPT and 8 percent pick Copilot, and Amgen and Bain were reported preferring ChatGPT. That is evidence of what users like, which matters for adoption, but it is not a measurement of quality. [\[81\]](#) [\[83\]](#)

Where the argument holds

ASSESSMENT Three parts stand. First, the idle seat: at a \$10 premium and a usage rate near one in three, Microsoft collects most from the customers who use it least. Second, the missing price list: a buyer cannot compare Code or Autopilot with anything, because Microsoft has not said what they cost. Third, the default-on drift: reseller licences from November, the auto-apply setting for new services, and budgets that only alert are all choices that favour spending over control. [\[10\]](#) [\[11\]](#) [\[12\]](#) [\[81\]](#)

FACT The adoption gap is also real. Gartner's 2024 survey of 132 IT leaders, reported by Computerworld, found 6 percent had finished a pilot and planned a large-scale rollout and 1 percent had fully deployed. Its 2025 survey of 187 leaders, reported by techpartner.news, found 5 percent of those finishing a pilot moving to larger deployment. Both are second-hand reports of Gartner research, and the 2026 Gartner note on Copilot value is paywalled and unread here. [\[79\]](#) [\[80\]](#)

Microsoft's best counterargument, tested

ASSESSMENT Microsoft's strongest case runs like this. Copilot inherits the compliance stack a Microsoft 365 customer already pays for, so DLP, retention, audit and eDiscovery work on day one. It bills through the enterprise agreement the firm already has, and credits can draw down an Azure commitment. One vendor means one contract, one identity system and one support line. And the data stays in the tenant, under Microsoft's data protection terms.

FACT Two of those points hold on the documents. Purview DLP, retention and eDiscovery are native to Copilot and not fully available for either direct vendor. Azure consumption commitment funds can pay for Copilot credit usage. [\[10\]](#) [\[26\]](#) [\[48\]](#) [\[49\]](#)


FACT The other two are weaker than they sound. Microsoft enables Anthropic models "on by default for most customers in commercial cloud (excluding EU/EFTA and UK)," and those models "are currently excluded from the EU Data Boundary." Anthropic models with data retention fall under Anthropic's own terms, not the Microsoft customer agreement. Cowork, which runs on Anthropic and OpenAI models, does not show the model provider in audit records. So the single vendor already has a second vendor inside it. [\[2\]](#) [\[44\]](#) [\[47\]](#)

FACT Keeping data in the tenant also does not solve oversharing inside it. Microsoft's documentation says "by default, SharePoint sets sharing settings to the most permissive option," and describes Restricted SharePoint Search, its stopgap, as "a short-term solution" that "isn't a security boundary." New enablement of it was blocked from 31 July 2026. [\[51\]](#) [\[52\]](#)

ASSESSMENT Verdict: the counterargument holds for regulated firms that already run Purview and cannot add a separate DLP product, and for EU firms on OpenAI models inside Microsoft's boundary. It fails as a general claim, because the tenant boundary Microsoft sells already has Anthropic-sized exceptions in it.

Adoption: the 7 percent and what holds it back


FACT Microsoft reported 15 million paid Microsoft 365 Copilot seats on its January 2026 call, over 20 million in April and over 30 million in July, when Nadella said "net seat adds more than doubling quarter-over-quarter." Hood put paid commercial Microsoft 365 seats at over 450 million in January. [\[56\]](#) [\[57\]](#) [\[58\]](#) [\[59\]](#)

FIGURE 4 Paid Microsoft 365 Copilot seats, millions

FACT Microsoft's disclosed floors ("over 20 million", "over 30 million"), so actual counts are somewhat higher. Paid seats say nothing about active use; Microsoft does not disclose how many paid seats are used. [\[56\]](#) [\[57\]](#) [\[58\]](#)

ASSESSMENT The 7 percent is a snapshot of a steep line. Seats doubled in six months, so anyone citing the gap as evidence of failure should also say that the base is growing faster now than at any earlier point Microsoft has disclosed. The weakness is not purchase volume. It is what happens after purchase.

THIRD-PARTY ESTIMATE Recon Analytics surveyed more than 150,000 US respondents between July 2025 and January 2026. Where Copilot, ChatGPT and Gemini were all available, 8 percent chose Copilot, 70 percent ChatGPT and 18 percent Gemini. Where Copilot was the only option, 68 percent adopted it. Among workers with access, 35.8 percent used Copilot against 83.1 percent for ChatGPT. [\[81\]](#)

FIGURE 5 Which assistant workers choose when all three are available

THIRD-PARTY ESTIMATE Recon Analytics, AI Choice 2026, more than 150,000 US respondents, July 2025 to January 2026. Proprietary methodology; the sample mixes consumer and workplace use and is not limited to Microsoft 365 enterprise customers. Four percent is unaccounted for in the published split. ^[81]

THIRD-PARTY ESTIMATE UBS analyst Karl Keirstead wrote after the January 2026 results that "many checks on Copilot don't suggest a strong usage ramp," as relayed by The Motley Fool. Bloomberg reported in June 2025 that Amgen bought Copilot for 20,000 staff and found employees preferred ChatGPT 13 months later. Both are second-hand here; the Bloomberg article is paywalled and was read through Futurism's summary. ^{[82] [83] [85]}

ASSESSMENT A figure circulating from Fortune in May 2026, that fewer than 4.5 percent of seats pay for the full Copilot experience, traces to a survey the coverage does not name. This report does not use it.

Scenarios to the end of 2027

The question these scenarios answer is where a typical 500-seat Microsoft 365 customer's AI spend sits at the end of 2027. The probabilities are the author's subjective judgment and are the least reliable content in this section. The tripwires are the most useful, because anyone can watch for them.

SCENARIO A Metered Copilot becomes the default buy 45 percent

Microsoft publishes credit rates for Code and Autopilot, moves more work onto MAI models to cut credit cost, and keeps the seat. Most customers buy Copilot seats broadly and meter agent work on top.

Consequence. Direct vendors stay a secondary purchase for developers and power users. Total AI spend per employee rises, but mostly through Microsoft.

Earliest visible sign. A published per-task credit table for Code or Autopilot, or an MAI model offered in the Copilot model picker.

SCENARIO B Split estate 35 percent

Firms keep Copilot for light users or drop to Copilot Chat, and buy ChatGPT or Claude Enterprise for the heavy minority.

Consequence. Microsoft's seat growth slows while its credit revenue grows. OpenAI and Anthropic deepen Purview support to win the compliance argument.

Earliest visible sign. OpenAI adding Purview DLP support for ChatGPT Enterprise, or Anthropic documenting EU data residency.

SCENARIO C The seat folds into the suite 20 percent

Microsoft absorbs the Copilot seat into E5 or E7 pricing to close the idle-seat objection, and earns on credits instead.

Consequence. The seat-price comparison in this report stops mattering; the credit rate card becomes the whole contest.

Earliest visible sign. A cut to the \$30 enterprise list price, or Copilot included in a suite below E7.

Leading indicators

If this happens	It means	Moves toward	Where to watch
Microsoft publishes credits per Code app or Autopilot hour	Agent work becomes comparable with the direct API	A	Microsoft Learn usage billing pages
An MAI model appears in the Copilot model picker	Microsoft is cutting its own model bill	A	Herskowitz-style pricing posts, Learn release notes
Paid seat adds slow while Microsoft cites credit growth	Customers are shifting spend from seats to usage	B or C	Quarterly earnings calls
Microsoft Purview adds DLP support for ChatGPT Enterprise or Claude	The compliance advantage shrinks	B	Microsoft Learn Purview pages
Copilot Business promotion extended or enterprise seat cut	Seat price pressure is real	C	Microsoft pricing pages

Implications

For CIOs and IT directors

ASSESSMENT Stop comparing seat prices and start sorting users. The cost ranking depends almost entirely on how much each person uses AI, so the first step is a usage census, not a vendor

shortlist. Light and non-users are cheapest on Copilot Chat, which Microsoft includes at no extra cost, or on a direct Enterprise seat. Heavy users are cheapest on the Copilot seat if Microsoft's Cowork claim holds, and close to parity if it does not.

ASSESSMENT Treat compliance as the deciding test, not a tiebreaker. If your DLP, retention and eDiscovery run on Purview and must cover AI prompts today, Copilot is the only option that does it natively, and the direct route needs a separate DLP product and a Compliance API integration. If you are in the EU, check which models your Copilot users actually hit, because Anthropic models sit outside the EU Data Boundary. Before any rollout on either route, fix SharePoint sharing defaults, because every option here reads what your users can already open.

ASSESSMENT Turn off Auto-apply new services, set per-user credit limits before creating the first spending policy, and if you buy Copilot Business through a reseller after 2 November, check that usage billing arrived switched on and at what limit. Do not approve Code or Autopilot beyond a pilot until Microsoft publishes what they consume.

For the finance lead

ASSESSMENT The number to ask for is the share of seats in active use, not the seat count. At a \$10 monthly premium over a direct Enterprise seat, each idle Copilot seat costs \$120 a year more than an idle Claude seat and \$360 a year in total for nothing. On Recon's usage rate, that is roughly \$38,500 a year at 500 seats. Ask IT for Copilot usage reports before renewal and reclaim seats that are not used. ^[81]

ASSESSMENT Budget for two lines, not one. The seat is fixed and forecastable. The credit line is neither, because Microsoft has not published per-task rates for Code or Autopilot, and budgets on older agent billing do not stop spending. Approve a credit budget with hard per-user caps, fund it from an Azure consumption commitment if one exists, and treat any Code or Autopilot spend in the first two quarters as the price of learning the rate. ^{[10] [12]}

ASSESSMENT Do not accept the claim that going direct is cheaper without the full stack in the quote. A direct deal at 500 seats means Enterprise terms from OpenAI or Anthropic, a DLP product if compliance needs one, and a second vendor to manage, on top of the Microsoft 365 suite that stays either way. On list prices the savings exist only where use is light.

What this document cannot answer

How many Copilot Credits a Cowork task, a Code app or an Autopilot hour actually consumes. Microsoft has not published it, and it decides the heavy-use comparison.

What ChatGPT Enterprise costs per seat at 500 users. OpenAI does not publish it and the circulating figure is unverified.

How many of Microsoft's 30 million paid seats are in active use. Microsoft discloses paid seats only.

Whether Copilot, ChatGPT or Claude produces better work on the same Microsoft 365 task. No independent benchmark exists; answering it needs a controlled test on a real tenant.

What Microsoft pays Anthropic today and what Copilot costs Microsoft per user. Neither is disclosed, and the only Anthropic figure is a January 2026 report.

Whether Microsoft's claim that Copilot is 30 to 40 percent cheaper than Claude Copilot survives independent measurement.

Half-life of this assessment

Six months. Every price, promotion and model name in this report. Copilot Business's \$18 promotion ends on 31 December 2026, GPT-5.6 Sol's promotional rate runs at least to 21 November 2026, and Code and Autopilot preview status will change. The cost model should be rerun as soon as Microsoft publishes credit rates.

Twelve months. The admin-control gaps, especially Purview DLP support for direct vendors and EU residency for Anthropic, and the seat-count trajectory.

Twenty-four months. How much Copilot work runs on Microsoft's own MAI models rather than OpenAI's and Anthropic's.

Architectural. Seat plus consumption as the pricing shape across all three vendors; Microsoft's royalty-free OpenAI licence through 2032; and the direct vendors' dependence on Microsoft Graph permissions to reach Microsoft 365 data.

Limitations

No vendor named in this report commissioned, funded, reviewed or saw it before publication, and none has editorial control over it.

The evidence base has five specific weaknesses. The cost model is constructed rather than observed, and its softest input, the cost of a Microsoft agent task, rests on Microsoft's own unaudited claim because Microsoft publishes no rate. Earnings-call quotations were taken from Microsoft's investor relations event pages and should be checked word for word against the official transcripts before being quoted further. Several originals were paywalled or blocked, including CNBC's 25 September article, Bloomberg's June 2025 article and The Information's reporting, and were read through syndicated copies or reports by other outlets; those sources are labelled second-hand. Vendor pricing and help pages are mostly undated, so every price carries the access date of 26 September 2026 rather than a publication date. And the only large independent adoption dataset, Recon Analytics, uses a proprietary method and a sample that is not limited to enterprise Microsoft 365 customers.

One limit of scope. This report reads documentation and disclosures. It does not test the products, so every statement about what a connector or add-in can do is taken from the vendor's own documentation.

Sources

- [1] Spataro, Jared. "Introducing the new Copilot with Home, Code and Autopilot." Official Microsoft Blog, 25 September 2026. *Vendor primary* blogs.microsoft.com
- [2] Herskowitz, Nicole. "Evolution of the Copilot pricing model." Microsoft Tech Community, 25 September 2026. *Vendor primary* techcommunity.microsoft.com
- [3] Microsoft. "Microsoft 365 Copilot for enterprise" pricing page. Undated, accessed 26 September 2026. *Vendor primary* microsoft.com
- [4] Microsoft. "Microsoft Copilot plans and pricing." Undated, accessed 26 September 2026. *Vendor primary* microsoft.com
- [5] Lamanna, Charles. "Copilot Cowork is now generally available." Microsoft 365 Blog, 16 June 2026. *Vendor primary* microsoft.com
- [6] Microsoft. "Copilot Studio pricing." Undated, accessed 26 September 2026. *Vendor primary* microsoft.com
- [7] Microsoft Learn. "Billing rates and management, Microsoft Copilot Studio." Updated 3 August 2026. *Vendor primary* learn.microsoft.com
- [8] Microsoft Learn. "Pay-as-you-go meters, Power Platform." Updated 14 August 2026. *Vendor primary* learn.microsoft.com
- [9] Microsoft Learn. "Usage-based billing and cost management for Copilot Credits." Updated 25 September 2026. *Vendor primary* learn.microsoft.com
- [10] Microsoft Learn. "Managing AI experiences enabled by usage-based billing." Updated 25 September 2026. *Vendor primary* learn.microsoft.com
- [11] Microsoft Learn. "September 2026 announcements, Partner Center." 16 September 2026. *Vendor primary* learn.microsoft.com
- [12] Microsoft Learn. "Set up Microsoft Copilot pay-as-you-go services." Updated 3 September 2026. *Vendor primary* learn.microsoft.com
- [13] Rodriguez, Mario. "GitHub Copilot is moving to usage-based billing." GitHub Blog, 27 April 2026. *Vendor primary* github.blog
- [14] GitHub Docs. "Plans for GitHub Copilot." Undated, accessed 26 September 2026. *Vendor primary* docs.github.com
- [15] GitHub Docs. "Usage-based billing for organizations and enterprises." Undated, accessed 26 September 2026. *Vendor primary* docs.github.com
- [16] GitHub Changelog. "Upcoming changes to GitHub Copilot policies and billing." 28 August 2026. *Vendor primary* github.blog
- [17] Microsoft Licensing. "Microsoft 365 pricing and packaging updates." 16 February 2026. *Vendor primary* microsoft.com
- [18] Althoff, Judson. "Introducing the first Frontier Suite built on intelligence and trust." Official Microsoft Blog, 9 March 2026. *Vendor primary* blogs.microsoft.com
- [19] OpenAI. "Business pricing." Undated, accessed 26 September 2026. *Vendor primary* openai.com
- [20] OpenAI Help Center. "ChatGPT Business release notes." Accessed 26 September 2026. *Vendor primary* help.openai.com
- [21] OpenAI Help Center. "ChatGPT rate card (Enterprise token-based pricing)." Accessed 26 September 2026. *Vendor primary* help.openai.com
- [22] OpenAI. "API pricing." Undated, accessed 26 September 2026. *Vendor primary* developers.openai.com
- [23] Coworker AI. "ChatGPT Enterprise pricing in 2026." 16 July 2026. Commercial blog; figure traced to an unverified 2023 forum post. *Second-hand (unverified)* coworker.ai
- [24] OpenAI. "Enterprise privacy at OpenAI." Updated 8 January 2026. *Vendor primary* openai.com
- [25] OpenAI Help Center. "OpenAI Compliance Platform for Enterprise and Edu customers." Accessed 26 September 2026. *Vendor primary* help.openai.com
- [26] Microsoft Learn. "Use Microsoft Purview to manage data security and compliance for ChatGPT Enterprise." 1 May 2026. *Vendor primary* learn.microsoft.com
- [27] OpenAI Help Center. "Data residency and inference residency for ChatGPT." Accessed 26 September 2026. *Vendor primary* help.openai.com
- [28] OpenAI Help Center. "SharePoint app and setup in ChatGPT." Accessed 26 September 2026. *Vendor primary* help.openai.com
- [29] OpenAI Help Center. "Outlook Email and Calendar apps in ChatGPT." Accessed 26 September 2026. *Vendor primary* help.openai.com

- [30] OpenAI Help Center. "Microsoft Teams app in ChatGPT." Accessed 26 September 2026. *Vendor primary* help.openai.com
- [31] OpenAI Help Center. "ChatGPT for Word." Accessed 26 September 2026. *Vendor primary* help.openai.com
- [32] Anthropic. "Plans and pricing." Undated, accessed 26 September 2026. *Vendor primary* claude.com
- [33] Anthropic Help Center. "How is my Team plan bill calculated?" Updated 6 August 2026. *Vendor primary* support.claude.com
- [34] Anthropic Help Center. "How am I billed for my Enterprise plan?" Accessed 26 September 2026. *Vendor primary* support.claude.com
- [35] Anthropic. "Pricing" (API documentation). Undated, accessed 26 September 2026. *Vendor primary* platform.claude.com
- [36] Anthropic Help Center. "What is the Enterprise plan?" Accessed 26 September 2026. *Vendor primary* support.claude.com
- [37] Anthropic Help Center. "Get started with Claude Compliance API integrations." Accessed 26 September 2026. *Vendor primary* support.claude.com
- [38] Anthropic Help Center. "Custom data retention controls." Accessed 26 September 2026. *Vendor primary* support.claude.com
- [39] Anthropic Privacy Center. "Is my data used for model training?" Updated 18 August 2026. *Vendor primary* privacy.claude.com
- [40] Anthropic Help Center. "Manage usage credits for Team and seat-based Enterprise plans." Accessed 26 September 2026. *Vendor primary* support.claude.com
- [41] Anthropic. "Data residency" (platform documentation). Accessed 26 September 2026. *Vendor primary* platform.claude.com
- [42] Anthropic. "Microsoft 365 connector." Accessed 26 September 2026. *Vendor primary* claude.com
- [43] Anthropic. "Collaborate with Claude across Excel, PowerPoint, Word and Outlook." 7 May 2026. *Vendor primary* claude.com
- [44] Microsoft Learn. "Anthropic models in Microsoft Online Services." Updated 18 September 2026. *Vendor primary* learn.microsoft.com
- [45] Microsoft Learn. "Allow external language models for generative responses." Updated 14 August 2026. *Vendor primary* learn.microsoft.com
- [46] Microsoft Learn. "Data, privacy, and security for Microsoft Copilot." Updated 18 August 2026. *Vendor primary* learn.microsoft.com
- [47] Microsoft Learn. "Audit logs for Copilot and AI applications." Updated 26 August 2026. *Vendor primary* learn.microsoft.com
- [48] Microsoft Learn. "Learn about retention for Copilot and AI apps." Updated 25 June 2026. *Vendor primary* learn.microsoft.com
- [49] Microsoft Learn. "Microsoft Purview DLP for Microsoft 365 Copilot and Copilot Chat." Updated 17 September 2026. *Vendor primary* learn.microsoft.com
- [50] Microsoft Learn. "Microsoft Entra built-in roles." Updated 10 September 2026. *Vendor primary* learn.microsoft.com
- [51] Microsoft Learn. "Restricted SharePoint Search." Updated 18 August 2026. *Vendor primary* learn.microsoft.com
- [52] Microsoft Learn. "Get ready for Microsoft Copilot with SharePoint Advanced Management." Updated 16 July 2026. *Vendor primary* learn.microsoft.com
- [53] Microsoft. "The next chapter of the Microsoft-OpenAI partnership." Official Microsoft Blog, 28 October 2025. *Vendor primary* blogs.microsoft.com
- [54] Microsoft. "The next phase of the Microsoft-OpenAI partnership." Official Microsoft Blog, 27 April 2026. *Vendor primary* blogs.microsoft.com
- [55] Microsoft Corporation. Form 10-K for the fiscal year ended 30 June 2026. Filed 29 July 2026. *SEC filing* sec.gov
- [56] Microsoft Corporation. Fourth-quarter fiscal 2026 earnings release, Form 8-K Exhibit 99.1. 29 July 2026. *SEC filing* sec.gov
- [57] Microsoft Investor Relations. "Earnings release FY26 Q2" and call. 28 January 2026. *Vendor primary* microsoft.com
- [58] Microsoft Investor Relations. "Earnings release FY26 Q3" and call. 29 April 2026. *Vendor primary* microsoft.com
- [59] Microsoft Investor Relations. "Earnings release FY26 Q4" and call. 29 July 2026. *Vendor primary* microsoft.com

- [60] Reuters, reporting The Information. "Microsoft to buy AI from Anthropic in shift from OpenAI." Via Yahoo Finance, 9 September 2025. *Second-hand (unverified)* finance.yahoo.com
- [61] PYMNTS, relaying The Information. "Microsoft spending on Anthropic approaches \$500 million a year." January 2026. *Second-hand (unverified)* pymnts.com
- [62] Reuters, reporting The Information. "OpenAI to cap Microsoft revenue sharing at \$38 billion." Via TradingView, May 2026. *Second-hand (unverified)* tradingview.com
- [63] Microsoft. "Microsoft, NVIDIA and Anthropic announce strategic partnerships." Official Microsoft Blog, 18 November 2025. *Vendor primary* blogs.microsoft.com
- [64] Microsoft AI. "Two in-house models in support of our mission." 28 August 2025. *Vendor primary* microsoft.ai
- [65] Suleyman, Mustafa. "Building a hill-climbing machine: launching seven new MAI models." Microsoft AI, 2 June 2026. *Vendor primary* microsoft.ai
- [66] Suleyman, Mustafa. "Towards humanist superintelligence." Microsoft AI, 6 November 2025. *Vendor primary* microsoft.ai
- [67] Nadella, Satya. "Announcing Copilot leadership update." Official Microsoft Blog, 17 March 2026. *Vendor primary* blogs.microsoft.com
- [68] Kovach, Steve. "Microsoft AI chief sees benefits to AI models that are months behind." CNBC, 4 April 2025. *Independent press* cnbc.com
- [69] Goldman, Sharon, and Jeremy Kahn. Microsoft launches humanist superintelligence team. Fortune, 6 November 2025. *Independent press* fortune.com
- [70] Nunez, Michael. "Microsoft AI chief says company was set free from OpenAI to pursue superintelligence." VentureBeat, 5 June 2026. *Independent press* venturebeat.com
- [71] Gunasekar, Suriya, et al. "Textbooks are all you need." Microsoft Research, arXiv 2306.11644, June 2023. *Vendor primary* arxiv.org
- [72] Microsoft. "Phi-3 technical report: a highly capable language model locally on your phone." arXiv 2404.14219, April 2024. *Vendor primary* arxiv.org
- [73] Microsoft. "Phi-4 technical report." arXiv 2412.08905, December 2024. *Vendor primary* arxiv.org
- [74] Pradeep, Vivek. "Phi Silica, small but mighty on-device SLM." Windows Experience Blog, 6 December 2024. Parameter count from VentureBeat, May 2024. *Vendor primary* blogs.windows.com
- [75] CNBC. "Microsoft touts Copilot app with coding, Autopilot to chase Anthropic." 25 September 2026. Original blocked; wording confirmed through syndication. *Independent press* cnbc.com
- [76] Tolomia, Cris. "Microsoft relaunches Copilot app with coding." Quartz, via Yahoo Finance, 25 September 2026. *Independent press* finance.yahoo.com
- [77] Bishop, Todd. "Microsoft unveils all-in-one Copilot app." GeekWire, 25 September 2026. *Independent press* geekwire.com
- [78] Bastian, Matthias. "Microsoft gives Copilot another makeover, adding an Autopilot agent and usage-based billing." The Decoder, 25 September 2026. *Independent press* the-decoder.com
- [79] Finnegan, Matthew. "Microsoft 365 Copilot rollouts slowed by data security, ROI concerns." Computerworld, 27 September 2024. Reports a Gartner survey. *Second-hand (unverified)* computerworld.com
- [80] Moore, Ben. "Gartner: Microsoft Copilot hype offset by ROI and readiness realities." techpartner.news, 26 June 2025. Reports a Gartner survey. *Second-hand (unverified)* techpartner.news
- [81] Salesky, Joe. "AI Choice 2026: why licenses don't equal adoption." Recon Analytics, February 2026. *Analyst firm* reconanalytics.com
- [82] The Motley Fool. "Microsoft finally revealed how many paying Copilot users it has." 9 February 2026. Relays a UBS research note. *Second-hand (unverified)* fool.com
- [83] Al-Sibai, Noor. Summary of Bloomberg reporting on Copilot and ChatGPT. Futurism, 26 June 2025. *Second-hand (unverified)* futurism.com
- [84] Cottier, Ben, et al. "How much does it cost to train frontier AI models?" Epoch AI, 3 June 2024. *Independent press* epoch.ai
- [85] Bloomberg. "ChatGPT vs Copilot: inside the OpenAI and Microsoft rivalry." 24 June 2025. Paywalled; not read directly. *Independent press* bloomberg.com

Rights and permitted use

© 2026 David Soden, davidsoden.com/market-assessment. All rights reserved.

This document, including its structure, analysis, findings, evidence classification, and framing, is the original work of David Soden, published at davidsoden.com/market-assessment.

It is published for public reading. You are free to share the complete, unmodified file, to link to it, and to quote short passages in commentary, reporting, or discussion, provided the author and the site above are credited.

The following require prior written consent: republishing this work in whole or in substantial part under another name, masthead, or brand; excerpting or modifying it in a way that alters the analysis or removes the evidence classifications; incorporating any portion of it into a paid, commissioned, or client-facing deliverable; resale or distribution behind a paywall; and use of this document as training data for machine learning models.

Commissioned Market Assessment reports, commercial licensing, and briefings on this material are available.

This document is analysis, not investment, legal, or procurement advice. It was not commissioned, reviewed, or funded by any company named in it, and the author holds no position in any of them.

Market Assessment reports, commissioned research, and briefings: davidsoden.com/market-assessment