David Soden

MARKET ASSESSMENT | AUGUST 2026

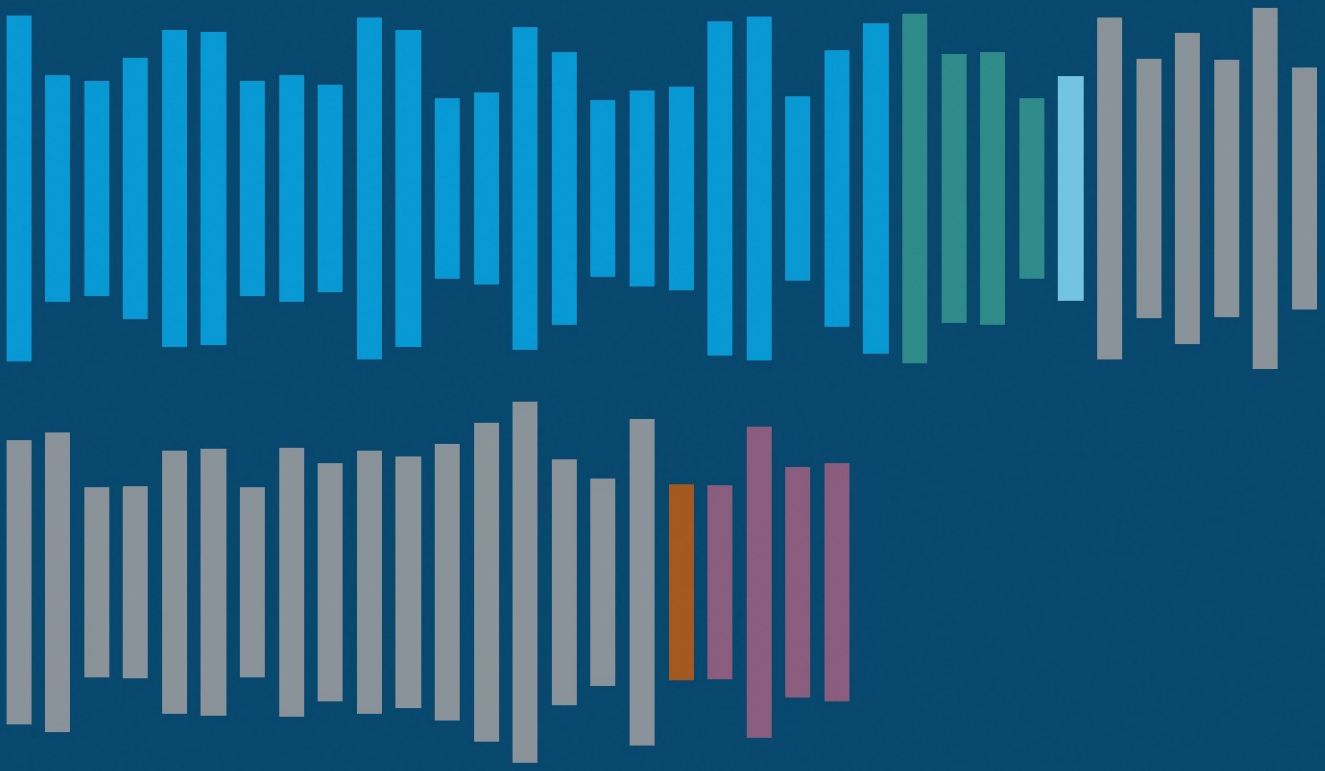

56 CLASSIFIED CLAIMS · FACT 23 VENDOR 4 ESTIMATE 1 ASSESS 23 HYPO 1 SCEN 4

Copilot Is Not One Product

Six commercial models behind one brand: what each Microsoft Copilot surface actually bills, where the boundaries sit, and how to budget a category that reprices twice a year

David Soden

Independent analyst and practitioner, enterprise automation and applied AI

davidsoden.com/market-assessment

REPORT DS-2026-6818 · 2026 · VERSION 1.0 · PUBLIC

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.

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 assesses the commercial structure of the Microsoft Copilot portfolio as of August 2026: what each Copilot-branded surface bills for, which meter it runs on, which administrative console governs it, which invoice it lands on, and what a buyer is actually committing to when they approve a Copilot rollout. It is written for the people who carry the consequences: the end user deciding whether to ask for a seat, the administrator and architect who must make the tenant safe enough to turn it on, and the finance owner who has to forecast the bill.

The subject is Microsoft, but the analysis is about pricing architecture rather than product quality. Nothing here argues that Copilot works badly. Several of the components are good. The argument is that the portfolio is commonly described as one integrated system and is commercially structured as six separate ones, and that the gap between those two descriptions is where budget overruns and stalled rollouts come from.

Published commentary on this category is unusually contaminated. Microsoft publishes seat counts without denominators and growth rates without bases. Partners and resellers publish licensing explainers while selling deployment services. Competing vendors publish cost breakdowns of Microsoft products on pages that sell an alternative. Comparison sites publish ranked pricing tables assembled from other comparison sites. All four appear in search results looking equally authoritative, which is why the same question returns four different numbers.

Every substantive claim below therefore carries one of six classifications, and any claim resting on an external source carries a numbered reference linking to the source list at the end. Where evidence is thin, second-hand, or commercially interested, that is stated in the sentence rather than hidden in the footnote. Prices are United States list rates in US dollars. Enterprise Agreement and CSP pricing is negotiated and undisclosed, so nothing here describes what any specific customer pays.

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

Four premises are assumed rather than demonstrated. Each paragraph names which conclusions fail if the premise is wrong.

One. Published list pricing reflects the shape of the commercial model

Negotiated enterprise pricing is invisible. This report assumes that discounting changes the magnitude of a bill but not its structure, so a customer paying half of list for Copilot Studio credits is still on a consumption meter with the same behavior at the margin. If Microsoft is routinely selling flat-rate all-inclusive agreements that suppress metering entirely, the cost-predictability argument in this report weakens considerably for large enterprises, though not for the mid-market.

Two. Microsoft's disclosed seat counts are accurate

Paid seat figures from earnings releases are assumed honest. They are unaudited in the sense that no third party verifies the definition of a paid seat, and Microsoft does not disclose how many are in multi-year commitments versus active use. The adoption analysis depends on these

figures. The commercial-structure analysis does not, and would survive intact if the seat numbers were materially overstated.

Three. Documentation reflects current billing behavior

Rate tables on Microsoft Learn are assumed to describe what customers are actually charged. Microsoft states that rates change and that the Learn page is the source of truth over any snapshot. Several rates cited here were verified in June and July 2026 by practitioners rather than by the author against a live invoice. If the published rates and the billed rates diverge, the specific credit arithmetic in this report is wrong even though the model shape is not.

Four. The metering direction of travel continues

The report assumes Microsoft keeps moving AI capability from seeded entitlements into consumption meters. Five separate moves in twelve months point this way. If Microsoft reverses and folds AI into base suite pricing the way Google did, Scenario B becomes the base case and most of the budgeting guidance here becomes conservative rather than wrong.

Executive summary

The Copilot portfolio is routinely diagrammed as a hub with spokes, one AI layer radiating across productivity, data, development, security, and business applications. That picture is wrong in a specific and expensive way. There is no hub. There are six commercial models with different units of consumption, different administrative consoles, different invoices, and different budget owners, joined only by a shared brand and a shared identity plane.

The practical consequence is that the person who approved the rollout usually did not approve the thing that generates the bill. A finance team that signs off on a per-seat number is not looking at the token meter, the compute-unit meter, the credit meter, or the capacity rental sitting underneath it. Microsoft has not hidden any of this. It is documented in public. It is simply distributed across enough surfaces that assembling it takes deliberate work, and the marketing does not encourage that work.

FINDING 1 **The portfolio is six meters, not one product.**

> **ASSESSMENT** Per-seat add-on, independent per-seat plus token consumption, rented capacity, hourly provisioned compute units, consumption credits, and a per-user governance layer. Each has a different unit, a different console, and a different bill. A buyer who models Copilot as a seat cost has modeled one of six. [1] [3] [4] [8] [13]

FINDING 2 **The base seat is a prerequisite, not the product.**

> **FACT** Microsoft 365 Copilot at $30 per user per month is an add-on that requires a qualifying Microsoft 365 or Office 365 license underneath it. Nothing in the Copilot brand can be bought as a standalone replacement for the suite, so the seat number quoted in most comparisons understates the committed spend by the entire cost of the base plan. [29] [25]

FINDING 3 **AI capacity is migrating out of licenses and into meters.**

> **FACT** AI Builder capacity add-ons went end of sale on 1 November 2025. All AI Builder credits, including the seeded capacity bundled into Power Platform and Dynamics licenses, are removed on 1 November 2026, with no conversion of entitlement into Copilot Credits and different rates on each side of the boundary. Capacity organizations already owned becomes capacity they must purchase. [2]

FINDING 4 GitHub Copilot did the same thing on the developer side.

> **FACT** On 1 June 2026 GitHub replaced premium request units with AI Credits, where one credit equals one cent and usage is metered against token consumption at each model's published rate. Seat prices did not change. What the seat includes did. Model multipliers, and with them the cheap-model fallback that softened quota exhaustion, were removed. [8] [9] [17]

FINDING 5 Credit consumption varies by more than an order of magnitude within one product.

> **FACT** Microsoft documents a classic answer at 1 credit, a generative answer at 2, an agent action at 5, and Microsoft Graph tenant grounding at 10. Microsoft's own example describes a single grounded complex prompt consuming 12 credits. The same agent, differently built, produces bills that differ by tenfold or more on identical traffic. [1] [10]

FINDING 6 Adoption after three years remains a single-digit percentage of the installed base.

> **FACT** Microsoft reported over 30 million paid Microsoft 365 Copilot seats at FY26 Q4 against a Microsoft 365 commercial base last disclosed as over 450 million and estimated at roughly 464 million. That is under 7 percent, from the vendor with the deepest enterprise distribution in software. Seat growth accelerated sharply in the quarter, which makes the trajectory interesting and the current penetration modest. [7] [18] [24]

FINDING 7 Google took the opposite commercial path, and that is the sharpest available contrast.

> **ASSESSMENT** Google retired its Gemini add-ons and folded AI into Workspace base plans at a higher plan price, effective for existing customers from March 2025. Microsoft kept AI as a priced add-on and then added meters beneath it. One vendor made AI a property of the seat. The other made it a property of usage. Buyers evaluating the two are not comparing like products. [16]

Category structure

The organizing insight is that Copilot is a naming system rather than a product, and the naming system spans six distinct commercial models. Below is the structure as a buyer meets it, with published United States list rates.

Commercial model	Representative product	Published rate (US list)	Administered and invoiced through
Per-seat add-on	Microsoft 365 Copilot	$30 per user per month, annual commitment, requires a qualifying base license	Microsoft 365 admin center; M365 invoice
Independent per-seat plus token meter	GitHub Copilot Business and Enterprise	$19 and $39 per user per month, plus AI Credits metered on token consumption since 1 June 2026	GitHub billing; separate agreement
Rented capacity	Copilot in Power BI and Fabric	Requires paid Fabric F2 or higher, or Power BI Premium P1 or higher; Pro and PPU do not qualify	Azure or Power BI capacity; Azure invoice
Hourly provisioned compute	Microsoft Security Copilot	Security Compute Units billed hourly on provisioned capacity, with a separate overage rate	Azure subscription; Azure invoice
Consumption credits	Copilot Studio, Copilot Cowork, AI Builder successor workloads	$200 per 25,000-credit pack per month, or roughly $0.01 per credit pay-as-you-go	Power Platform admin center and Azure; split invoices
Per-user governance layer	Microsoft Agent 365	$15 per user per month, GA 1 May 2026, or bundled in Microsoft 365 E7 at $99	Microsoft 365 admin center; M365 invoice

ASSESSMENT Read down the right-hand column and the structural problem appears. The same agent can produce a line on the Microsoft 365 invoice for its governance seat and a line on the Azure invoice for its execution, owned by two teams with different budgets and no shared view. The Power BI equivalent is older and better understood: per-seat Pro licensing on the Microsoft 365 side, capacity consumption on the Azure side. The agent era reproduces that split rather than resolving it. [13] [20]

FIGURE 1 **Five seat prices, four invoices, one brand**

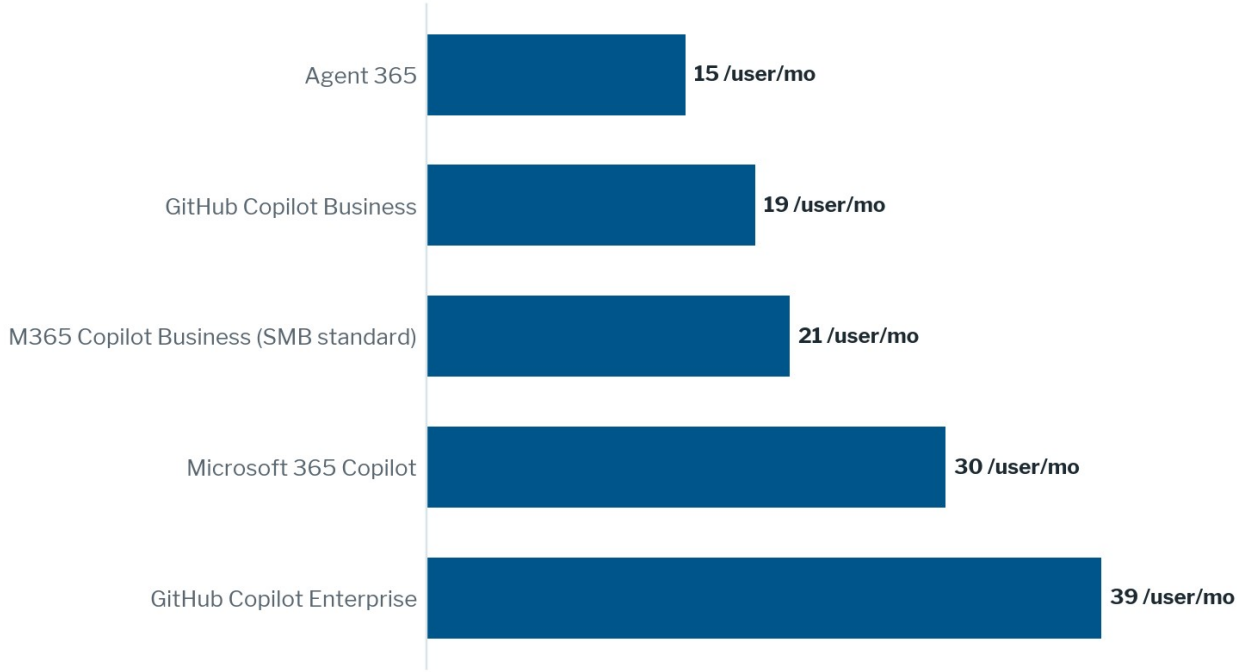

FACT Published US list rates, per user per month. Every one of these sits on top of a base license that is not shown, and none of them covers the consumption meters underneath. The SMB rate has carried a promotional price whose expiry date is reported differently by different sources. [8] [13] [25] [29]

FACT One meter has a genuine zero-rated boundary and it matters. When a user holding a Microsoft 365 Copilot license interacts with an agent inside Copilot Chat, Microsoft Teams, or SharePoint, classic answers, generative answers, and Microsoft Graph tenant grounding do not draw down Copilot Studio capacity, subject to fair use limits. Everything outside that boundary meters: external-facing agents, unlicensed users, and standalone channels. [1] [11]

ASSESSMENT That boundary is the single most useful thing an architect can hold in their head, and it is also the reason cost forecasts break. The boundary is defined by three things at once, the user's license, the surface the agent runs on, and the feature invoked. A design change on any one of them moves an agent from free to metered without anything visibly changing for the user. [1] [11]

Sizing, and why to discount it

VENDOR CLAIM Microsoft reported more than 30 million paid Microsoft 365 Copilot seats for the quarter ending 30 June 2026, with net seat additions more than doubling quarter over quarter, and said the number of customers with more than 50,000 seats grew more than sevenfold year over year. [7]

FACT The comparable base figure is older. Microsoft last disclosed Microsoft 365 commercial paid seats as over 450 million in January 2026 and did not refresh the number at FY26 Q4, reporting only 6 percent annual seat growth. Applying that growth rate gives roughly 464 million, which puts Copilot penetration under 7 percent. [7] [18]

FIGURE 2 **The numerator Microsoft publishes, against the denominator it does not**

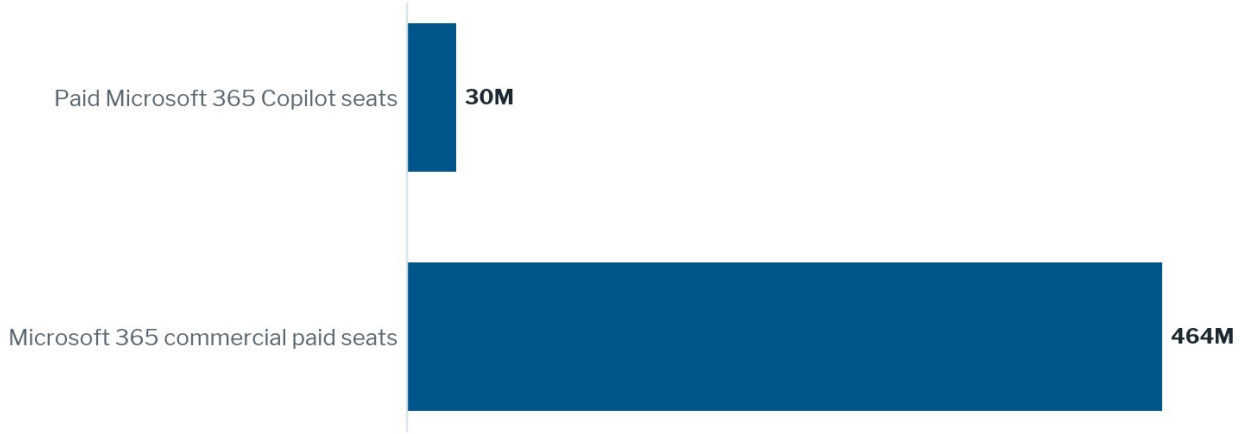

ASSESSMENT Seats in millions. The Copilot figure is Microsoft's own disclosure for the quarter ending 30 June 2026. The base is not a disclosure: it is the January 2026 figure of over 450 million carried forward at the 6 percent seat growth Microsoft reported, so treat it as an estimate with the arithmetic shown rather than a reported number. [7] [18]

ASSESSMENT The growth rate and the penetration rate tell opposite stories and both are true. Tripling paid seats in two quarters is a real commercial result. Reaching under 7 percent of your own installed base after nearly three years, with unmatched distribution, is a real adoption problem. Anyone citing only one of these numbers is arguing rather than analyzing.

VENDOR CLAIM A widely circulated figure places Microsoft's AI business at a $37 billion annual run rate growing 123 percent year over year. That figure belongs to the FY26 Q3 disclosure in April 2026 and was not refreshed at Q4. It also covers the whole AI business, not Copilot. Treating it as a current Copilot number, as several summaries do, is wrong on both the date and the scope. [7] [18]

ASSESSMENT Growth stated as a percentage without a disclosed base is close to uninformative, and 123 percent growth on an undisclosed mix of Azure AI consumption, Copilot seats, and model serving tells a buyer nothing about whether their own Copilot deployment will pay back. The useful vendor number in this category is the seat count with a denominator, and Microsoft supplies the numerator far more readily than the denominator.

Reference vendor: what Microsoft has published

Facts first, drawn from Microsoft and GitHub primary documentation. Vendor characterizations follow separately.

FACT Microsoft 365 Copilot is $30 per user per month on annual commitment and operates only as an add-on to a qualifying base subscription. Microsoft's own deployment documentation treats the qualifying base license as a prerequisite for the seat to activate at all. [29]

FACT A small-business variant sits below it at $21 per user per month standard list, having run at a promotional $18. The expiry of that promotion is reported as 31 March 2026, 30 June 2026, 30 September 2026, and 31 December 2026 by four different licensing sources, all of them writing after at least one of those dates had passed. The disagreement is itself the evidence: nobody

outside Microsoft can state the current promotional terms with confidence, which is a poor foundation for a budget. [25] [26] [27]

FACT Copilot Studio has no per-seat price. Capacity is bought as a prepaid pack of 25,000 Copilot Credits for $200 per tenant per month, or pay-as-you-go at roughly $0.01 per credit through an Azure subscription. Microsoft states that unused credits do not carry over to the next month and that credit usage resets on the first day of each month regardless of when the license was acquired. [1] [11]

Copilot Studio feature	Copilot Credits	Pay-as-you-go cost at $0.01
Classic (authored) answer	1	$0.01
Generative answer	2	$0.02
Agent action (trigger, connector call, topic transition)	5	$0.05
Microsoft Graph tenant grounding	10	$0.10
Microsoft's documented example: one grounded complex prompt	12	$0.12

FIGURE 3 **The same agent, built two ways, costs tenfold apart**

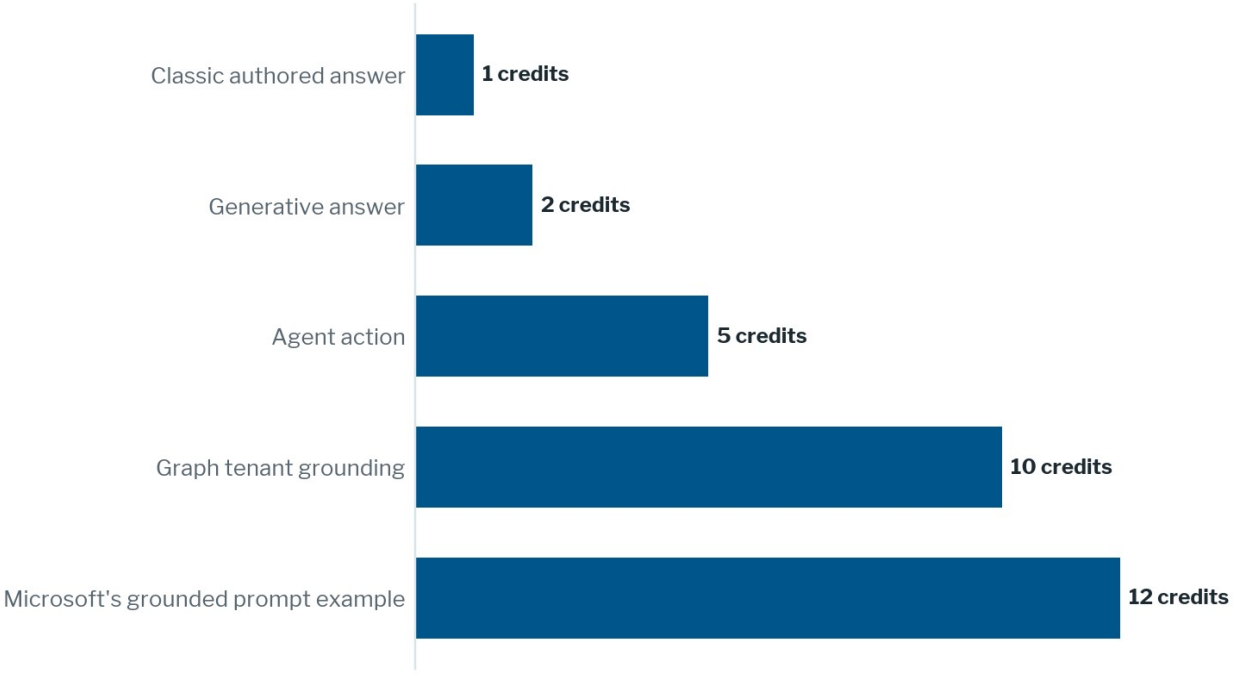

FACT Copilot Credits per event, from Microsoft's published rate card. These are unit rates, not observed averages: no public data exists on how a real agent's traffic distributes across these events, which is why a credit forecast built from this card is a design estimate rather than a measurement. [1]

ASSESSMENT The spread on this card is the whole cost story. Grounding a generative answer on the tenant takes it from 2 credits to 10, a fivefold step, and the gap from a scripted answer to a grounded one is tenfold. Tenant grounding is the feature that makes an agent useful on company data and it is the most expensive rate on the card. An organization that follows the obvious design instinct, ground everything on the tenant, has chosen the highest unit cost available before a single user types anything. [1]

FACT Exceeding purchased capacity is not merely a billing event. Microsoft documents that Copilot Studio enforces purchased capacity monthly, that unused credits do not carry over, and that enforcement can result in service denial rather than an overage charge unless pay-as-you-go is separately enabled. [1] [11]

FACT On the developer side the seat and the meter are now separate line items on the same invoice. GitHub Copilot Business is $19 and Enterprise $39 per user per month. Since 1 June 2026 every plan carries a monthly allotment of GitHub AI Credits at one cent per credit, pooled at the billing-entity level for Business and Enterprise, consumed by chat, agents, code review, and CLI usage but not by code completions or next edit suggestions, with additional usage purchasable beyond the allotment. [8] [9]

FACT GitHub documents that existing Business and Enterprise customers received a raised allowance for the first three months of usage-based billing, from 1 June to 1 September 2026, after which included usage returns to the standard amounts. It also documents that when pooled credits are exhausted and additional usage is not allowed, Copilot is blocked until the next cycle. [9]

FACT Microsoft Security Copilot bills through provisioned Security Compute Units charged hourly, with a minimum of one provisioned SCU and overage capacity billed on usage. Microsoft's pricing page states $4 per provisioned SCU and works an example in which 3.5 SCUs consumed inside an hour against 4 provisioned units bills as $16 for that hour. Billing is calculated in whole clock hours, so provisioning at 9:05 and again at 9:45 bills two units inside the same hour. [4] [5]

FACT Microsoft has made an SCU allocation available to eligible Microsoft 365 E5 and E7 tenants at 400 SCUs per month for every 1,000 user licenses, capped at 10,000 SCUs per month, on a consumption-deducted pool rather than hourly provisioned capacity. [4] [6]

THIRD-PARTY ESTIMATE The overage rate of $6 per SCU is consistent across independent licensing practitioners but is not stated on the Microsoft pricing page in the sources reviewed here, and at least one practitioner source says overage billing against the E5 inclusion pool was not yet available and that exceeding the allocation results in throttling. Treat the model as reliable and the overage number as needing confirmation against a quote. [19]

FACT Copilot in Power BI requires a paid Fabric capacity at F2 or higher, or Power BI Premium at P1 or higher. Microsoft's documentation states plainly that a Power BI Pro or Premium Per User license alone is not sufficient. Trial capacities, free SKUs, and workspaces in Embedded license mode are not supported. [3] [12]

FACT Microsoft Agent 365 became generally available on 1 May 2026 at $15 per user per month, or bundled inside Microsoft 365 E7 at $99 per user per month alongside E5, Microsoft 365 Copilot, and the Entra Suite. The standalone license is per user rather than per agent and requires a qualifying base of E5, Defender and Purview Suite frontline worker, or Business Premium. [13] [20]

FIGURE 4 What E7 bundles, and what it does not

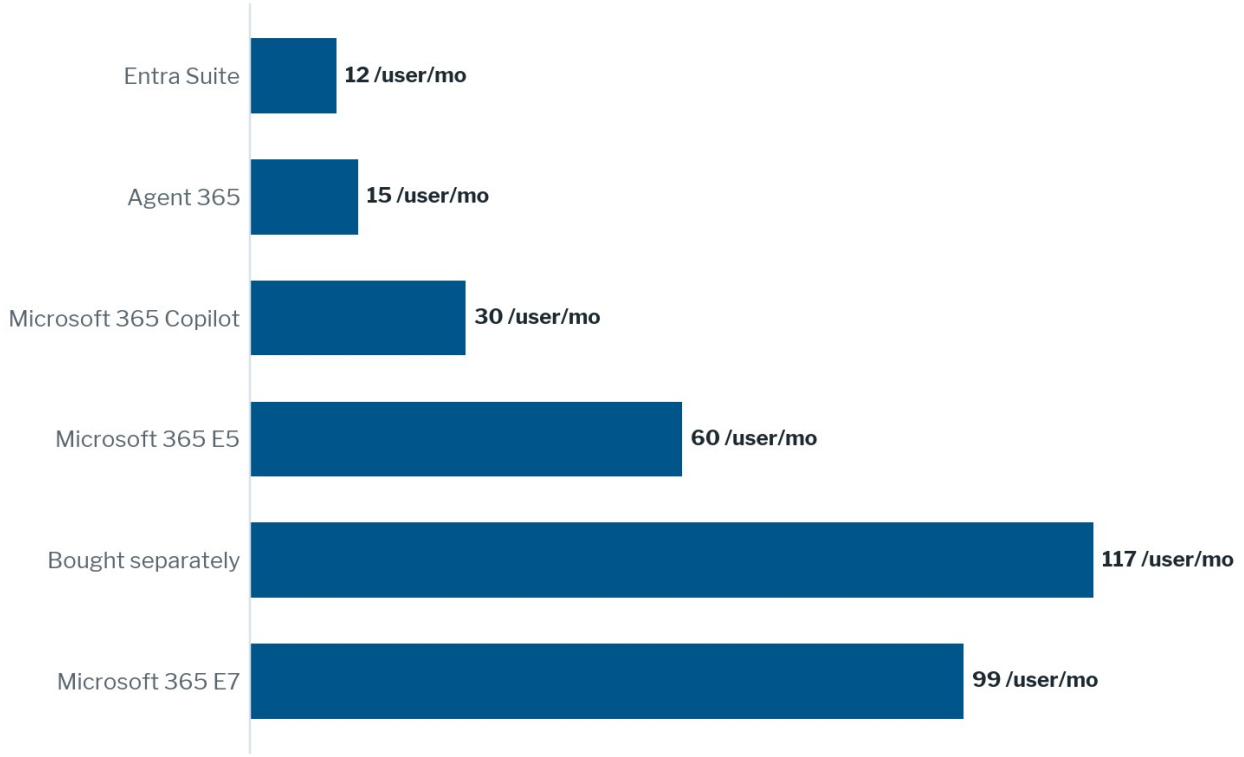

FACT US list rates per user per month. The 15 percent saving is real at list. What the bundle does not cover is the consumption underneath it: Copilot Credits for agents built in Copilot Studio, Azure tokens for agents built in Foundry, and SCUs beyond the included pool. [13] [21] [28]

What Microsoft claims

VENDOR CLAIM Satya Nadella framed the FY26 Q4 seat figure as reflecting the confidence customers place in Microsoft to power their AI transformation, and described Copilot as evolving from chat to Cowork to Autopilots. [7]

VENDOR CLAIM Microsoft positions the credit model as a single common currency for agents across products, and describes the September 2025 rename from messages to Copilot Credits as carrying no change in pack quantity or pay-as-you-go rate. [1]

ASSESSMENT The common-currency claim is true within Copilot Studio and misleading across the portfolio. Copilot Credits are one currency. GitHub AI Credits are a different currency on a different meter with its own allotments and its own overage. Security Compute Units are a third. Fabric capacity units are a fourth. Describing any of these as unified is accurate only inside the boundary of the product that issued them. [1] [4] [8]

Reading the 2025 to 2026 release wave

Seven commercial changes in roughly twelve months. Read together they describe a deliberate migration from seeded entitlement to metered consumption, with one genuine consolidation running the other way.

Change	Strategic reading	What would confirm it
Messages renamed to Copilot Credits, 1 September 2025	Establishes a portable consumption currency before the products that will consume it exist	Further Microsoft products adopting Copilot Credits as their billing unit rather than inventing new ones
Role-based Dynamics Copilots folded into the $30 M365 Copilot seat	Genuine simplification where a separate SKU was suppressing attach rate rather than earning revenue	Continued retirement of standalone role Copilot SKUs, with capability moving into the base seat
AI Builder add-ons end of sale, 1 November 2025	Closes the old capacity currency to new customers before the replacement is mandatory	No renewal or true-up path offered after November 2026, which Microsoft has already documented
Dynamics 365 Premium SKUs seeded with 1,000 Copilot Credits per user	Uses credits as an upsell lever from standard to premium tiers rather than as a pure meter	Standard-tier SKUs continuing to receive no seeded credits while premium allocations grow
Agent 365 GA at $15 per user, 1 May 2026	Prices agent governance as a per-seat product, separate from agent execution	Governance features remaining outside the E5 and Entra Suite bundles rather than being absorbed
Microsoft 365 E7 at $99 per user, 1 May 2026	Rebundles the expensive add-ons into a single higher tier to lift average revenue per seat	A second Frontier-tier SKU, or E7 attach becoming the primary Copilot sales motion
GitHub AI Credits replace premium requests, 1 June 2026	Moves developer AI from flat subscription to token metering while holding seat prices flat	Included allowances shrinking at promo expiry, which GitHub has already scheduled for 1 September 2026

FACT The Dynamics consolidation is real and worth stating clearly, because it cuts against this report's own argument. The standalone $50 per user Copilot for Sales and Service SKUs were retired and their capability moved into the $30 Microsoft 365 Copilot license. A customer previously paying separately for both is now paying less. [22] [23]

ASSESSMENT That consolidation and the AI Builder retirement happened in the same portfolio within a year, in opposite directions. The pattern that reconciles them is not simplification or complication but margin: Microsoft consolidates where a separate SKU was suppressing adoption of a seat it wants attached, and unbundles where usage can be metered. Both moves are rational. Neither is evidence of an integration strategy. [2] [22]

FACT The Dynamics credit seeding illustrates the same logic. Since November 2025 the four Premium SKUs each include 1,000 Copilot Credits per user per month pooled at tenant level, while the standard Enterprise and Professional SKUs receive none and must buy credits separately. [23]

Competitive assessment

The comparison set is smaller than the product count suggests, because most Copilot-branded surfaces compete with different vendors.

Vendor	Position	Principal constraint
Microsoft	Only vendor with AI inside the incumbent productivity suite, the developer platform, the security stack, and the business applications	Six meters, four consoles, and no unified cost view; the buyer assembles the total themselves
Google (Workspace)	AI bundled into base plans since the add-on retirement, at a raised plan price	Suite displacement is a far harder sale than an add-on; wins on pricing clarity, loses on incumbency
Google (Gemini Enterprise)	Separate agent and enterprise-search platform priced per seat with consumption beyond quota	Reintroduces exactly the metering complexity that bundled Workspace Gemini avoids
OpenAI	Strong standalone assistant and API with no productivity-suite dependency	No document, mail, or identity substrate; competes for usage, not for the tenant
Anthropic	Strong enterprise assistant and coding position; supplies model capability inside Microsoft's own Copilot surfaces	Same substrate gap as OpenAI; presence inside a competitor's product is reach and dependence at once
Salesforce, ServiceNow	Agent platforms anchored in their own systems of record with their own consumption units	Compete with Copilot Studio and Dynamics agents specifically, not with the productivity seat
Cursor and coding-tool entrants	Direct competitors to GitHub Copilot on developer workflow and model choice	Compete with one of the six meters and are irrelevant to the other five

ASSESSMENT Two entries on that list are routinely mispositioned. Anthropic and OpenAI are described as Copilot competitors in most comparisons, but they compete for a knowledge worker's attention rather than for a tenant's identity, permissions, and document estate. A company can and frequently does buy both. Conversely, Cursor is described as a Copilot competitor and that is precisely correct, but only against GitHub Copilot, which shares a brand with Microsoft 365 Copilot and nothing else.

ASSESSMENT The Google contrast is the most instructive because both vendors solved the same problem and chose opposite answers. Google raised base plan prices by a few dollars a seat and gave AI to every user, which makes the AI invisible in procurement and removes the attach conversation at renewal. Microsoft kept AI as a decision, which preserves the ability to charge power users more and guarantees a recurring negotiation. Google's approach is better for buyers who want predictability. Microsoft's is better for buyers who genuinely only need AI for a subset of staff, and there are more of those than the marketing on either side admits. [16]

What actually decides purchases

The criteria that appear on evaluation checklists are largely not the criteria that decide whether a Copilot deployment proceeds, expands, or stalls.

ASSESSMENT The first real gate is tenant permission hygiene, and it is a gate rather than a criterion because it can stop a rollout outright. Copilot does not bypass permissions. It makes a decade of accumulated oversharing searchable in plain English, which converts a dormant governance debt into an immediate incident. Organizations discover this in pilot, with executives, who hold the broadest access in the company. [14] [15]

ASSESSMENT The second is whether finance can forecast the bill within a tolerable band. A per-seat number is approvable because it is a headcount multiplication. A consumption meter whose unit cost varies tenfold with agent design is not approvable by the same process, and in most organizations there is no established owner for a variable AI line item that appears on the Azure invoice rather than the software one. [1]

ASSESSMENT The third is the base license prerequisite, which decides more deals than it appears to. Copilot cannot be bought as a standalone. For an organization sitting on lower-tier plans, evaluating Copilot means evaluating a suite upgrade, and the conversation changes from a $30 decision to a full renegotiation. [29] [25]

ASSESSMENT The checklist criteria that decide comparatively little: model quality benchmarks, context window comparisons, feature-by-feature matrices against Gemini, and the count of applications a Copilot appears in. These are the things vendor materials compete on and the things buyers argue about before discovering that the permissions audit will take four to eight weeks.

Failure modes, ranked by commercial impact

Ordered by how often each one ends or freezes a purchase, which is not the order in which they are complained about.

One. The consumption bill nobody owns

ASSESSMENT The dominant failure is organizational rather than technical. The seat is approved by IT or HR, the credits are consumed by a business unit that built an agent, and the invoice arrives on a cloud cost centre managed by a third team. Because agent design sets the unit cost before traffic does, the same nominal deployment produces wildly different bills, and nobody in the chain owns the variable. Microsoft ships capacity monitoring and per-environment allocation in the Power Platform admin center precisely because tenant-pooled credits are otherwise consumed with no attribution. [10]

Two. Permission debt surfacing in pilot

ASSESSMENT Second most common and the most likely to produce an outright stop. The pattern is consistent: a leadership pilot surfaces a document from a forgotten site, no policy is breached, and Legal or HR freezes the rollout pending cleanup. Microsoft ships Purview and SharePoint Advanced Management to address this, includes SAM with the Copilot license, and publishes a

staged oversharing blueprint structured as pilot, deploy, operate. That is a reasonable response and also an admission that the tenant is usually not ready. [14] [15]

Three. Capacity gating on the analytics side

ASSESSMENT Teams assume Power BI Copilot comes with their Pro seats, discover it requires F2 or P1 capacity, and either abandon the use case or absorb an unbudgeted capacity rental. This is a smaller financial failure than the first two but a more frequent source of the feeling that the ecosystem does not fit together, because the two products carry the same brand and the same vendor and still require a separate purchase. [3] [12]

Four. Service denial at capacity exhaustion

FACT Because credits do not roll over and enforcement can deny service, an agent that runs out of tenant capacity mid-month stops working. GitHub's pooled model has the same shape on the developer side: when pooled AI Credits are exhausted and additional usage is not permitted, Copilot is blocked until the cycle refreshes. When the exhausted capacity was consumed by an unrelated department, the failure lands on a team with no visibility into the cause. [1] [9] [11]

Five. Activation without adoption

ASSESSMENT Seats assigned and not used. This is the failure most often measured and least often decisive, because it produces waste rather than a stop, and waste is renewed once before it is cut. Independent commentary on the low paid-adoption rate consistently names naming confusion, unclear return, and compliance concern as contributing causes. [24]

Where value accrues

ASSESSMENT The structural force in this category is not model capability, which is converging and increasingly commoditized across vendors. It is metering. Whoever owns the unit of account owns the pricing power, and Microsoft has spent twelve months establishing Copilot Credits as a currency that spans Copilot Studio, AI Builder's successor workloads, Cowork, and parts of Dynamics. Establishing a currency before the workloads that consume it exist is the same move a cloud provider makes with compute units, and it works for the same reason. [1] [2]

ASSESSMENT The counterintuitive read is that the metering shift favors buyers with disciplined architecture and punishes everyone else, which inverts the usual enterprise software pattern. Under per-seat pricing, a badly designed deployment and a well designed one cost the same. Under credit metering, a team that grounds selectively and scopes agents narrowly pays a fraction of what an equivalent team pays for grounding everything by default. Cost control moves from procurement to design, and most organizations have no process that puts those two functions in the same room.

HYPOTHESIS Proposition offered for testing: the identity plane, not the model or the assistant, is where durable margin sits, which is why Agent 365 is priced per user at $15 and sold as a layer above E5 rather than inside it. Confirmed if agent governance stays outside existing bundles through 2027 while execution pricing continues to fall. Refuted if Microsoft absorbs Agent 365

into E5 to drive attach, which would signal it sees governance as a feature rather than a market. [13] [20]

Forecasting a category that resets

ASSESSMENT Point forecasts do not survive contact with this category, and the evidence is the subject's own rate of change. Within roughly twelve months the agent currency was renamed, a capacity currency was scheduled for retirement with no conversion path, the developer product's entire billing model was replaced, two new SKUs launched, role-based SKUs were retired into a different seat, and base suite prices rose. Any number attached to 2028 in this market is decoration. What follows is scenarios with subjective probabilities and, more usefully, observable events anyone can monitor without private information. The probabilities are the least reliable content in this report. The tripwires are the most useful.

Scenarios to end 2027

SCENARIO A Consumption becomes the primary meter 45 percent

> Seats settle into an entitlement floor and the commercially interesting revenue moves to credits, tokens, and compute units across the whole portfolio. Seeded allocations shrink at each renewal in the way AI Builder's did.
>
> **Consequence.** Budget ownership fractures further. Finance functions that cannot forecast AI spend either cap it hard, which suppresses usage, or stop trying, which produces the overruns that generate the FinOps tooling market.
>
> **Earliest visible sign.** A third Microsoft product retiring a seeded entitlement in favor of Copilot Credits, or included GitHub AI Credit allowances falling again after the September 2026 promotional expiry.

SCENARIO B Bundling reasserts under competitive pressure 30 percent

> The E7 pattern generalizes. Microsoft folds AI into higher base tiers rather than selling add-ons, converging on Google's model, and the standalone Copilot add-on becomes a legacy motion for customers who did not upgrade.
>
> **Consequence.** Pricing becomes predictable and materially higher at the base. Buyers lose the option to license AI for a subset of staff, which is worse for organizations with a small genuine power-user population.
>
> **Earliest visible sign.** A second Frontier-tier SKU, or Microsoft ending new sales of the standalone $30 add-on in favor of bundle attach.

SCENARIO C Buyer-forced simplification 25 percent

> Procurement pushback, aided by the visible gap against Google's bundled pricing, forces Microsoft to publish unified cross-product cost tooling and commit-based agreements that flatten the meters into a single negotiated envelope.
>
> **Consequence.** The commercial complexity described in this report becomes a transitional artifact rather than a structural feature, and the advantage shifts back to incumbency.

> **Earliest visible sign.** Microsoft shipping a single tenant-level view of Copilot Credits, GitHub AI Credits, SCUs, and Fabric capacity in one console, which does not currently exist.

FIGURE 5 **Three ways this ends, with the caveat that the numbers are judgment**

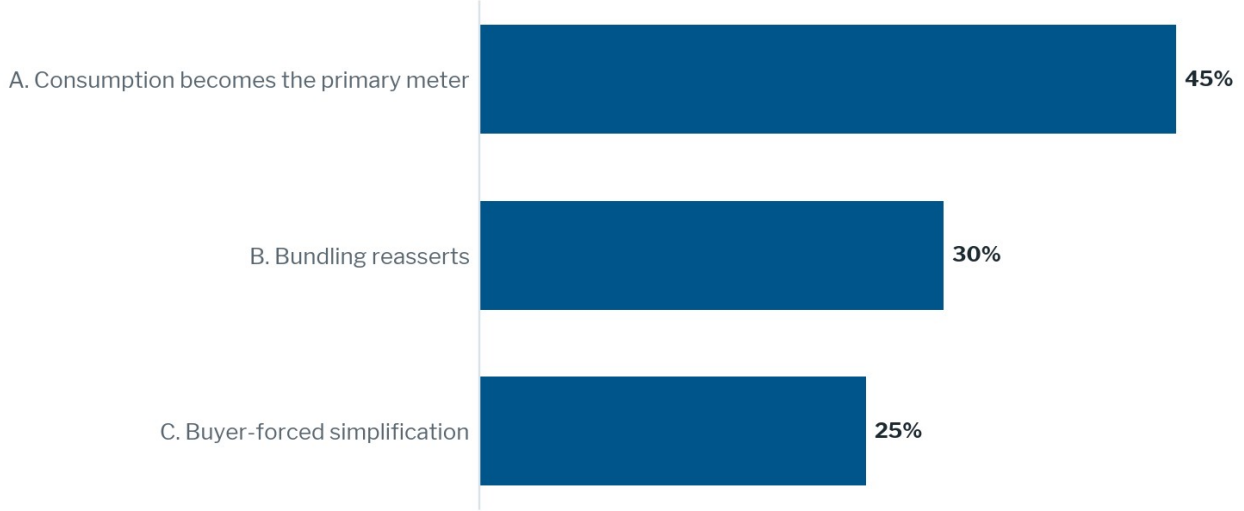

SCENARIO Subjective probabilities to end 2027, assigned by the author and summing to 100. They are not calculated, not derived from a model, and not the part of this section to rely on. The leading indicators in the next table are observable; these numbers are not.

Leading indicators

Observable without private information. Each maps to a scenario.

If this happens	It means	Moves toward	Where to watch
Seeded AI Builder credits are removed on schedule in November 2026	Microsoft is willing to withdraw entitlements customers already hold	A	Microsoft Learn AI Builder retirement page
GitHub Business allowances drop from promotional levels after September 2026	Included usage is the lever, not the seat price	A	GitHub billing documentation
A further Microsoft product adopts Copilot Credits as its billing unit	The common currency strategy is working	A	Product licensing guides and Learn rate pages
Agent 365 is absorbed into E5 or the Entra Suite	Governance is being treated as a feature, not a market	B	Microsoft 365 licensing announcements
A second Frontier-tier suite launches above or beside E7	Bundling is the primary revenue motion	B	Microsoft 365 roadmap and pricing pages

If this happens	It means	Moves toward	Where to watch
The standalone $30 Copilot add-on stops being sold to new customers	Add-on economics have lost to bundle economics	B	Partner and CSP price lists
Microsoft ships one console showing all four meters together	The cost-visibility complaint has reached product planning	C	Microsoft 365 and Power Platform admin centers
Commit-based Copilot agreements appear in standard CSP terms, not just EA	Metering is being flattened for mid-market buyers	C	CSP partner pricing announcements
Microsoft discloses a refreshed Microsoft 365 commercial seat denominator	Confidence that the penetration ratio now reads well	Any	Quarterly earnings releases
Copilot penetration passes 10 percent of the commercial base	Adoption has broken out of the early-majority stall	Any	Earnings releases and independent seat-math analyses
Google reintroduces a priced AI add-on above bundled Workspace	The bundled model is not covering inference cost	B	Google Workspace pricing pages
Independent FinOps tooling for Copilot Credits reaches meaningful adoption	The metering problem is durable enough to sustain a vendor category	A	Funding announcements and practitioner community discussion

Dated predictions

Stated so the document can be scored later.

Prediction	Resolution date	Confidence
Seeded AI Builder credits are removed on 1 November 2026 with no conversion into Copilot Credits, as documented	1 December 2026	High, roughly 90 percent
GitHub Copilot Business included AI Credits fall to the standard allowance after the promotional period ends, without a compensating seat price cut	31 December 2026	High, roughly 85 percent
Microsoft does not ship a single tenant-level console covering Copilot Credits, GitHub AI Credits, Security Compute Units, and Fabric capacity together	30 June 2027	Moderate, roughly 70 percent
Paid Microsoft 365 Copilot seats remain below 15 percent of disclosed Microsoft 365 commercial paid seats	30 June 2027	Moderate, roughly 65 percent
At least one additional Microsoft product moves from a seeded AI entitlement to Copilot Credit metering	31 December 2027	Moderate, roughly 60 percent

Implications

For end users considering a seat

The $30 seat is the part of the portfolio with the most predictable value and the least surprising cost, and for most individual knowledge workers it is also the only part that matters. The seat covers Copilot in the Office applications and zero-rates internal agent use inside Copilot Chat, Teams, and SharePoint. If someone is describing an integrated ecosystem to you, ask which of the six meters they mean, and ask who in your organization is watching the ones you would be spending against.

For administrators and architects

Do the permissions work before the license work, because the sequence is not reversible once a pilot surfaces something sensitive. Then set per-environment credit allocation in the Power Platform admin center before the first agent ships, treat tenant graph grounding as an opt-in decision per agent rather than a default, and establish which invoice each planned workload will appear on and who reads it. The single highest-leverage architectural choice available is whether an agent grounds on the tenant graph, because that decision alone moves the unit cost fivefold.

For finance and business owners

Budget three lines, not one: the base suite, the Copilot seats, and a variable consumption line that will land on the Azure invoice rather than the software one. Ask your Microsoft representative for the credit consumption estimate on your specific planned agents rather than a per-seat quote, and ask specifically what happens when capacity is exhausted, because the answer is service denial rather than an overage charge unless pay-as-you-go is separately enabled. If a vendor comparison shows Copilot at $30 against a competitor at a similar number, the comparison has almost certainly omitted the base license prerequisite on the Microsoft side.

For investors

The seat count is the wrong metric to track and the one Microsoft supplies most readily. The metric with information in it is whether consumption revenue grows faster than seat revenue, because that separates a durable platform position from an add-on attach cycle. Watch for the denominator: Microsoft's willingness to refresh the Microsoft 365 commercial seat number is itself a signal about how the penetration ratio currently reads.

What this document cannot answer

Four questions require primary research that public sources do not support. What large enterprises actually pay after Enterprise Agreement discounting, and whether negotiated terms suppress metering behavior at scale. What the real distribution of credit consumption per agent looks like across a population of production deployments, rather than the illustrative rates on the published card. How many assigned Copilot seats are in active weekly use, as distinct from paid. And whether the organizations that completed a permissions remediation before rollout show measurably different outcomes than those that did not, which is the single most useful number nobody has published.

Half-life of this assessment

The specific rates decay fastest. Credit rates, seat prices, promotional allowances, and SCU inclusion terms should be treated as accurate for roughly six months and re-verified against Microsoft Learn and the vendor pricing pages before any budget decision. The four conflicting expiry dates reported for one promotional price are a fair guide to how quickly this layer moves.

The release-wave reading and the scenario probabilities hold for around twelve months, subject to the leading indicators above. The structural claim, that the portfolio is six commercial models joined by a brand and an identity plane rather than one integrated product, is architectural and will survive several repricing cycles. It fails only if Scenario B runs to completion and AI becomes a property of the base seat across the whole portfolio.

Limitations

This assessment is independent. The author holds no vendor sponsorship, no reseller relationship, and no commercial arrangement with any organization named in it, and received no compensation from any party for its production or distribution.

Three specific weaknesses in the evidence base should be weighed. First, the pricing detail rests heavily on Microsoft's own documentation, which is authoritative on list rates and silent on realized pricing, so every cost statement here describes published behavior rather than any customer's invoice. Second, several figures, particularly the Security Compute Unit overage rate and the current status of promotional pricing, are corroborated across independent licensing practitioners rather than quoted from a live vendor pricing page, and practitioner sources on licensing frequently carry a consulting interest in complexity. Third, adoption and failure-mode analysis draws on partner and practitioner accounts that are self-selected toward organizations

that encountered problems, which will overstate frequency even where the mechanism described is real.

Sources

[1] Microsoft. "Copilot Studio billing and licensing." Microsoft Learn, updated 2026. *Vendor primary* learn.microsoft.com

[2] Microsoft. "End of AI Builder credits." Microsoft Learn, updated May 2026. *Vendor primary* learn.microsoft.com

[3] Microsoft. "Copilot for Power BI overview." Microsoft Learn. *Vendor primary* learn.microsoft.com

[4] Microsoft. "Microsoft Security Copilot pricing." Microsoft Security. *Vendor primary* microsoft.com

[5] Microsoft. "Security Compute Units and capacity." Microsoft Learn, April 2026. *Vendor primary* learn.microsoft.com

[6] Microsoft. "Learn about Security Copilot for Microsoft 365 E5 and E7 customers." Microsoft Learn, June 2026. *Vendor primary* learn.microsoft.com

[7] Microsoft. FY26 Q4 earnings press release and conference call, 29 July 2026. *Vendor primary* microsoft.com/investor

[8] GitHub. "GitHub Copilot is moving to usage-based billing." The GitHub Blog, 27 April 2026. *Vendor primary* github.blog

[9] GitHub. "Usage-based billing for organizations and enterprises." GitHub Docs. *Vendor primary* docs.github.com

[10] Microsoft. "Manage Copilot Studio credits and capacity." Microsoft Learn, Power Platform admin documentation. *Vendor primary* learn.microsoft.com

[11] Microsoft. "FAQ for Copilot Studio billing and licensing." Microsoft Learn, May 2026. *Vendor primary* learn.microsoft.com

[12] Microsoft. "Fabric Copilot capacity for usage billing." Microsoft Learn, June 2026. *Vendor primary* learn.microsoft.com

[13] Microsoft. "Microsoft 365 E7 and Agent 365 are now generally available." Microsoft Community Hub, 1 May 2026. *Vendor primary* techcommunity.microsoft.com

[14] Microsoft. "Mitigate oversharing to govern Microsoft 365 Copilot and agents." Microsoft Community Hub, September 2025. *Vendor primary* techcommunity.microsoft.com

[15] Microsoft. "Address oversharing in Microsoft 365 Copilot: deployment blueprint." Microsoft Learn. *Vendor primary* learn.microsoft.com

[16] Google. "The future of AI-powered work for every business." Google Workspace Blog, 15 January 2025. *Vendor primary* workspace.google.com

[17] Directions on Microsoft. "GitHub Copilot to move to usage-based pricing in June." 28 April 2026. *Analyst firm* directionsonmicrosoft.com

[18] Redmond, Tony. "FY26 Q4 Microsoft results sees Azure top $100 billion." Office 365 for IT Pros, 30 July 2026. *Independent press* office365itpros.com

[19] SAMexpert. "Security Copilot licensing and SCU pricing guide." 2026. *Practitioner review (consulting interest)* samexpert.com

[20] SAMexpert. "Agent 365 licensing: what it covers and costs." 12 June 2026. *Practitioner review (consulting interest)* samexpert.com

[21] SAMexpert. "Microsoft 365 E7: $99 bundle breakdown." 30 June 2026. *Practitioner review (consulting interest)* samexpert.com

[22] Stoneridge Software. "Big changes to Copilot licenses impacting Dynamics 365." 14 October 2025. *Third-party comparison (partner, commercial interest)* stoneridgesoftware.com

[23] arades GmbH. "Dynamics 365 licensing update May 2026: the 16 changes." 8 May 2026. *Practitioner review (consulting interest)* arades.de

[24] No Jitter. "4 obstacles impede paid Microsoft 365 Copilot adoption." 9 March 2026. *Independent press* nojitter.com

[25] Velosio. "Microsoft 365 Copilot pricing calculator." 8 July 2026. States the SMB promotional price runs through 31 December 2026. *Third-party comparison (partner, commercial interest)* velosio.com

[26] Copilot Experts. "Microsoft Copilot pricing explained: plans, cost and licensing." 27 June 2026. States the same promotion runs through 30 September 2026. *Third-party comparison (commercial interest)* copilot-experts.com

[27] Office Consumer. "How much is a Microsoft 365 Copilot license?" 26 April 2026. States the same promotion ran through 31 March 2026. *Second-hand (unverified)* officeconsumer.com

[28] TeamsFox. "Hidden costs IT teams miss in Microsoft Agent 365 pricing." 7 May 2026. *Practitioner review (commercial interest)* teamsfox.com

[29] Microsoft. "Set up Microsoft 365 Copilot and assign licenses." Microsoft Learn, May 2026. *Vendor primary* learn.microsoft.com

Rights and permitted use