Four credits, no exchange rate

Three of the four vendors now print a dollar price for an AI credit, and two print the same one: $0.01. Salesforce lists $500 per 100,000 and draws 20 for an action. Microsoft's rate card states no dollar figure. The price has converged and the exchange rate has not: how many credits a piece of work costs is still the vendor's to set, and the tenant pool ends chargeback.

3 November 2026 · Annika Sorensen, Writer, Strategic Intelligence Division · Operations

KEY POINTS

- The per-seat baseline no longer forecasts agent spend: three of the four vendors publish a dollar price for the credit and two publish one cent, so the divergence has moved off the price and onto the exchange rate, how many credits each draws for an action it defines itself.
- Microsoft prices an identical agent action at 5 Copilot Credits for a user without a Microsoft 365 Copilot licence and at no charge for a user with one, so the seat stays the product and the credit is a toll on everyone outside it.
- Credits pool at tenant or organisation level at Microsoft, Atlassian and Salesforce, so the vendor's invoice no longer names the team that spent the money and chargeback becomes an allocation the customer invents.
- Microsoft blocks an environment's agents at 125 per cent of prepaid capacity, including an agent in a conversation already under way, so an overrun arrives as an outage before it arrives as an invoice.
- Atlassian took the other side of the trade and filed it: providing AI tools at no or low cost to most Cloud customers may further increase hosting costs without corresponding revenue increases.

One GitHub AI Credit	One Salesforce Flex Credit, at list	One agent action, Microsoft	The same action, Microsoft 365 Copilot licence held
$0.01	**$0.005**	**5 credits**	**No charge**

GitHub Docs, usage-based billing for individuals; Salesforce Agentforce pricing page; Microsoft Learn, Copilot Studio billing rates. The Salesforce figure is the $500 per 100,000 credits printed on that page; Mosaic Singularity analysis, November 2026.

The price of an AI credit has narrowed to a cent or less, and the metering has not. GitHub converts an AI Credit at $0.01 in the document that defines it. Atlassian publishes the same one cent for usage above an allowance it includes inside the Jira, Confluence or Collection seat. Salesforce lists $500 per 100,000 Flex Credits, half a cent, and draws 20 of them for an agent action, ten cents at list for a unit it defines as one function. Microsoft's rate card prices an agent action at 5 Copilot Credits, states no dollar figure, and charges nothing where the user holds a Microsoft 365 Copilot licence. A CFO building an FY27 software budget on the FY26 per-seat baseline is now comparing four credit systems whose only common term is the word credit, and whose pooling rules decide whether any of it can be charged back.

The old unit ends at renewal, not on a date you choose

GitHub put its change on 1 June 2026 and gave Copilot Pro and Copilot Pro+ annual subscribers still on legacy billing three options. One of them is to stay, and it carries the ending: “When your annual plan ends, you'll be automatically downgraded to Copilot Free.” The request unit is not grandfathered. It expires with the term, and the customer who does nothing loses the paid plan as well as the unit. Source: GitHub Docs, “What changed with Copilot billing (legacy)”.

The unit that finance could forecast is gone

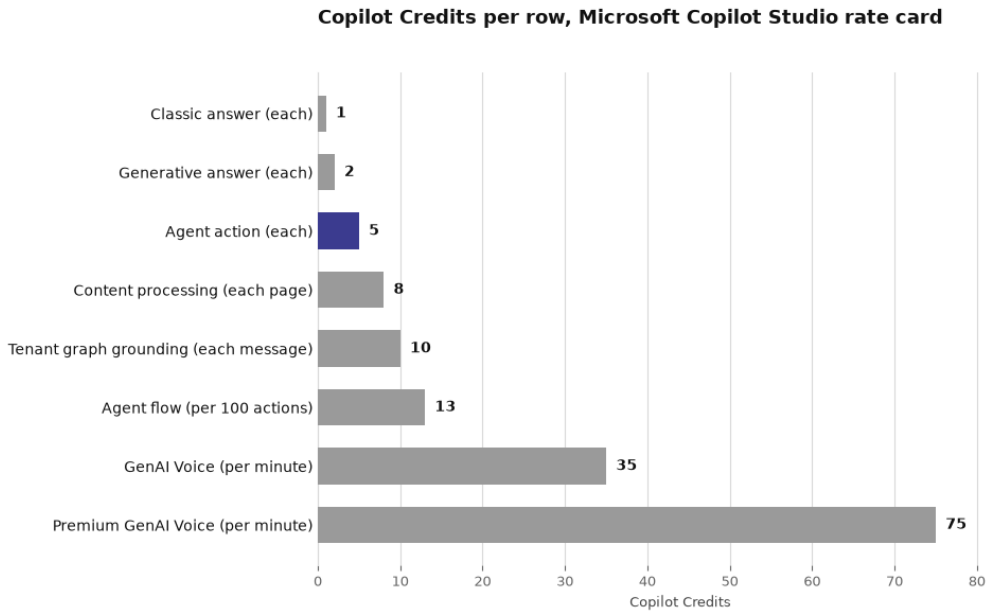
“Before June 1, 2026, billing was premium request-based: each model interaction cost one premium request unit (PRU), and a multiplier was applied based on which model you used” That was the shape of the old unit: a published table, a known unit, one price, and a forecast a controller could sign. From 1 June 2026 GitHub replaced it with tokens, and its documentation is explicit about the consequence: “the cost of an interaction therefore depends on two things: the model used and the number of tokens consumed” Tokens are produced by the work rather than chosen by the buyer, so the same task costs different amounts as a conversation lengthens or an agent loops. The allowance softens that without restoring the forecast. Copilot Pro includes 1,500 credits a month for $10, Pro+ 7,000 for $39, Max 20,000 for $100, and the balance resets at 00:00:00 UTC on the first day of the calendar month rather than on the subscription's billing date.

Microsoft made the same move a year earlier at no change in price. From 1 September 2025 its unit for agents changed from messages to Copilot Credits, and its licensing documentation records that there was “no change in the quantity per prepaid pack or to the pay-as-you-go rate”. The unit was renamed, not repriced. What did change is the vocabulary a buyer has to reconcile at renewal: a contract that said messages in 2025 is quoted in credits in 2027, and the Copilot Studio rate card states no dollar value for a credit at all, only the number of credits each row consumes.

Who sets the unit, and who only pays it

Salesforce keeps the multiplier and moves it out of the price. Its pricing page gives one number per action, 20 Flex Credits, and 30 for a voice action, then hands the buyer a second document: “Usage types and multipliers differ for Agentforce Voice actions.” The same page prices a hundred thousand credits at $500, which is half a cent a credit at list. The word action does not mean the same thing at the two vendors. Salesforce defines an action as “a specific function that an AI agent executes on the platform, such as updating a record, summarizing a complex case, answering a product inquiry, or executing a custom prompt or flow”. Microsoft defines agent actions as “steps such as triggers, deep reasoning, and topic transitions that appear on the activity map”, and prices a classic answer at 1 credit and tenant graph grounding at 10. Five credits and 20 credits are not a price comparison. They are two vendors metering different grains of work and both calling the unit an action.

Inside one rate card: Copilot Credits per row, on the basis each row is charged



Microsoft Learn, Copilot Studio billing rates, updated 31 October 2026. The rows are charged on different bases, shown in each label, so the bars are not a like-for-like comparison across rows. The agent action row is the row that matters for any comparison with Salesforce.

The dearest row on that card is 75 Copilot Credits and the cheapest is 1: a minute of Premium GenAI Voice against a classic answer. The bases differ, which is why the chart labels every row, and why a usage dashboard that totals credits tells a buyer almost nothing. What the card does fix is the failure mode. When a tenant reaches 125 per cent of prepaid capacity, enforcement blocks the environment's agents, including an agent in a conversation already under way, and the user is told the agent “has reached its usage limit”. An overrun is an outage before it is an invoice, and the threshold is the vendor's to enforce.

Who carries the cost when the unit is a pool

Four rate cards, one word, and what each does when the balance runs out

Vendor	Unit and where it is stated	Published price	What the allowance pools against	What happens when it runs out
GitHub	GitHub AI Credit, usage-based billing documentation	$0.01 per credit; 1,500 to 20,000 credits a month by plan	The plan, with an optional budget for extra usage	Unused credits forfeited at month end; a legacy annual plan downgrades to Copilot Free at term end
Salesforce	Flex Credit, Agentforce pricing page	$500 per 100,000 credits; 20 credits an action, 30 for a voice action	Across teams, channels and use cases	Not stated on the pricing page; voice multipliers sit in a separate rate card
Atlassian	Rovo credit, licensing page	$0.01 per credit over the allowance ($10.00 per 1,000); 25 to 1,500 credits a user a month included in the seat	A single organisation-wide pool	Extra usage billed at the current purchase price to an admin-set spending cap; switched off, new credit-consuming features pause until the allowance resets
Microsoft	Copilot Credit, Copilot Studio billing rates	5 credits an agent action; no dollar value stated	Across the entire tenant, with per-environment allocation optional	Agents blocked at 125 per cent of prepaid capacity

GitHub Docs, usage-based billing for individuals; Salesforce Agentforce pricing page; Atlassian Rovo licensing page; Microsoft Learn, Copilot Studio billing rates and licensing.

Three of the four pool the allowance above the seat, and each says so in its own words. Microsoft: “You pool this capacity across the entire tenant.” Atlassian: “All of an organization's allowances are combined into a single shared pool. Power users can consume more while light users consume less, with the balance spread across the org.” Salesforce: Flex Credits “can be applied across teams, channels, and use cases”. Pooling is a real benefit for adoption and it removes the control finance reaches for first. The invoice stops naming a team. Microsoft lets capacity be reserved per environment, and anything unallocated draws from the tenant, so a chargeback line exists only where someone has built it and keeps reconciling it.

> *“Our strategy to provide our AI tools at no or low cost to the majority of our Cloud customers may further increase our hosting costs without corresponding revenue increases.”*
>
> Atlassian, Form 10-K for the fiscal year ended 30 June 2026, filed 14 August 2026

That sentence is the trade in plain sight. Atlassian sells the seat, carries the inference to an allowance inside it, and prices only what runs past that allowance: $0.01 a credit, $10.00 per 1,000, and its own worked example turns a 28,000-credit allowance with 35,000 credits used into 7,000 extra credits and $70 for the month. The allowance is 25 credits a user a month on Jira Standard, 70 on Premium and 150 on Enterprise, and 250, 700 and 1,500 on the Collections. An allowance carried inside a seat price can be withdrawn from inside a seat price at the next renewal, without a rate card moving. A buyer who chose Atlassian because AI was included has bought a term, not a price.

The balance transfers nothing at renewal

None of this is an asset at the end of the term. GitHub forfeits unused credits at the month end and resets the balance on the first day of the calendar month. Atlassian's allowances refresh monthly on the usage billing date. Neither design lets a buyer bank agent capacity and carry it into a renewal negotiation. A seat at least told the buyer what a year would cost: seats, times price, times twelve. A credit balance produces a run rate, expires on the vendor's calendar, and reprices when the vendor changes a row. That is the difference between a capability that compounds inside a company and one rented by the month, and it is settled in the billing terms rather than in the product.

Who gains while the unit is unsettled

Microsoft's rate card is the clearest statement of who the credit is for. The billing table carries a second column headed “Used by Microsoft 365 Copilot licensed user”, and every row in it reads “No charge”. For employee-facing work, the same agent action costs a colleague without the licence 5 Copilot Credits and a colleague with it nothing. In Microsoft's own worked example, a tenant-grounded agent serving 50 licensed and 100 unlicensed users a day bills 4,800 Copilot Credits, and the arithmetic multiplies by the 100 unlicensed users only. Two footnotes in that document carry the scope a buyer should read before relying on it: the inclusion applies to employee-facing use in Microsoft 365 services and apps, it is subject to fair usage limits, and Computer-Using Agents are excluded. The seat is still the product. The credit is the toll on everyone outside it.

GitHub's conversion is the cleanest of the three that state a price, and it is worth saying why. One AI credit is one cent, in the same document that defines the unit, so a bid can be turned into dollars before signature and set against a second bid. Atlassian states the same one cent, but only for usage above an allowance priced inside the seat, so the buyer converts the marginal credit and never the included

one. Salesforce states half a cent and draws 20 credits for an action, ten cents a unit at list, with the multipliers for voice held in a separate rate card. Microsoft's rate card states the credits and no dollar figure, so capacity converts to dollars at the purchase and not in the document a budget is built from. That is the ask procurement should carry into every renewal in the FY27 cycle: state the dollar value of the unit and the basis it is charged on, or the contract prices nothing that can be forecast.

What each seat owes before 31 December 2026

- CFO: re-base the FY27 software budget line by line onto credits, and require two numbers in every AI renewal: the dollar value of a credit and the number of credits each unit of work draws. Three of these four vendors print a price and two print the same one, so the price per credit will not separate the bids; the credits each draws for a unit of work will. Microsoft states no dollar figure, so that request is either answered or refused in writing. By 31 December 2026.
- CIO: before signing a credit schedule, get the overage term in writing, meaning what happens at 125 per cent of prepaid capacity, whether agents are blocked, and who can lift the block. Microsoft publishes all three answers, which is what makes the question fair to put to every other bidder. By 31 December 2026.
- Procurement: put the four rate cards side by side per action and require each bidder to state the basis of its unit. An action is a single function at Salesforce and a step in an activity map at Microsoft, so a comparison that ignores the definition compares two products at once. By 31 December 2026.
- Financial controller: build the allocation before the renewal, not after it. Credits pool at tenant or organisation level at three of the four, so the invoice will not name the team that spent the money and the chargeback model has to stand without it. By 31 December 2026.
- CHRO: count the employee-facing agent users who hold no Microsoft 365 Copilot licence, because that population is where the metered cost sits while the licensed population is included within fair usage limits. By 31 December 2026.

What a board should put to management

The decision that reaches the board in this cycle is one of scope, not of spending: whether agent spend is held in a central technology line or charged back to the teams that incur it. The board resolves it by 31 December 2026, on three artefacts the seats above owe in the same window: the CFO's re-based FY27 budget, the controller's allocation built without the vendor naming a team, and the CIO's written overage term at 125 per cent of prepaid capacity. Asked before that scope is fixed, the five questions below return opinions; asked after it, they return the number the FY27 plan is built on.

- Which of our AI contracts price in credits, and what dollar value does each contract state for its unit?
- Does the contract fix the allowance, or can the vendor change what a seat includes at renewal?
- Which rows on the rate card does our workload actually hit, and which of them is the dearest?
- What happens to production agents at 125 per cent of capacity, and who is told before it happens?
- Can we evidence per-team agent usage from the vendor's data, or only from an allocation we build ourselves?

The test that settles it by 30 June 2027

- 31 December 2026: the FY27 renewals will have been signed or deferred by this date, and a credit contract either carries a dollar value per unit by then or does not.

- 1 June 2027: one year after GitHub's change. If the legacy annual plans have lapsed into Copilot Free and the one-cent conversion still stands, the fixed-dollar credit is a design choice rather than a transition.
- August 2027: Atlassian's next Form 10-K, on the pattern of the one filed on 14 August 2026. The sentence about providing AI tools at no or low cost either survives, with hosting costs rising against it, or it disappears.
- The price is settled at a cent or less wherever one is published, and the unit is not. If two of the four publish a common unit of work with a fixed dollar value before 30 June 2027, the divergence was a transition and a per-seat baseline can be rebuilt around it. If none does, the credit stays the budget unit and the exchange rate stays the vendor's to set.

Primary documents

- GitHub Docs, “What changed with Copilot billing (legacy)”, the change effective 1 June 2026.
- GitHub Docs, “Usage-based billing for individuals”, no publication date on the page; it sets out the usage-based billing that replaced request-based billing on 1 June 2026.
- Salesforce, Agentforce pricing page, dated 19 October 2026 in its page metadata, against a published date of 29 October 2024.
- Atlassian, Rovo licensing page, no publication date on the page; quoted for its plan allowances, read with the Form 10-K filed 14 August 2026.
- Microsoft Learn, “Billing rates and management”, Copilot Studio, document updated 31 October 2026.
- Microsoft Learn, Copilot Studio billing and licensing, last updated 15 October 2026.
- Atlassian Corporation, Form 10-K for the fiscal year ended 30 June 2026, filed 14 August 2026.