CURLY'S CONSULTING / RESEARCH PAPER 01

2026 Trends in AI monetization

Ammanuel "Tom" Selameab, CPP

Certified Pricing Professional | Curly's Consulting

November 2026

A comparative analysis of 20 documented company pricing changes and the buyer questions they leave unanswered.

Abstract

AI monetization is changing more than software price levels. It is changing what the customer buys: access, an included allowance, model consumption, a completed action, or a verified result. This paper examines 20 documented company pricing and packaging transitions that inform the 2026 market. The evidence covers changes originating from 2024 through November 7, 2026; historical origins and announced future billing are identified explicitly. The cases show several approaches coexisting: AI bundled into core subscriptions, differentiated tiers, seat-plus-credit hybrids, compute-sensitive charging, and outcome definitions that narrow or broaden the billable event. Public announcements establish what suppliers offered, but do not establish why a particular buyer will pay, how usage changes under the tariff, or whether the change caused better commercial results. The paper therefore treats competitive evidence as a starting point for pricing research. Before copying a model, a company should research its own buyers' workflow value, alternatives, package choices, comprehension of the bill, and willingness to accept usage or performance risk. Those observations, joined to delivery-cost data, are required to turn an observed market pattern into a defensible pricing decision. The conclusion sets out the research brief an investor or operator should commission.

1 Research scope and method

The unit of analysis is a documented transition in an existing software offer's price, package, charging metric, allowance, or purchase structure. Twenty company cases were selected because public primary sources explain an AI-linked commercial change. The selection includes infrastructure-adjacent tools, productivity suites, work platforms, and customer-service products. It is a purposeful comparative sample, not an estimate of market prevalence.

The evidence window ends November 7, 2026. Earlier transitions are retained where they explain the architectures visible in the 2026 debate. An event dated 2025 is not presented as a new 2026 change. A future effective date is not treated as an implemented charge. Where a support document states applicability but no original publication date, that distinction is preserved.

Each case records the prior arrangement, the revised arrangement, the relevant date and customer boundary, and the commercial question the transition raises. Dollar amounts are quoted in the units and billing context supported by the source. They are historical examples, not a current purchasing guide. Credits, actions, requests, resolutions, and tokens are not normalized as though they were equivalent units.

AI is an explicit part of the commercial redesign in the selected sources. That does not mean every change was caused solely by AI cost, or that other capabilities contributed no value. Announcements and billing documentation generally do not identify the causal effect on revenue, retention, or margin. The interpretations and proposed research questions below are our analysis of the documented choices.

2 The 20-company transition ledger

Company and event	Documented commercial transition
Microsoft, Dec. 2025	SMB-specific Copilot price and convenience bundles alongside the existing enterprise-oriented offer.
Google Workspace, 2025	Separate Gemini purchase becomes AI included in revised suite subscriptions.
Adobe, June/Aug. 2025	All Apps evolves into AI-rich Pro, with a limited-AI Standard alternative.
Notion, May 2025	Optional AI add-on becomes part of Business and Enterprise for the revised offer.
Atlassian, 2025	Previously separate Rovo becomes included in eligible paid cloud subscriptions.
Airtable, June 2025	Per-seat AI add-on ends; included credits and optional capacity packs remain.
Grammarly, Mar. 2025 onward	Eligible Premium subscriptions migrate to broader Pro with more AI prompts at the same bill.
Box, June 2024	Specified included AI-question caps become unlimited interactive questions in Enterprise Plus.
Asana, July/Aug. 2025	Included Basic access expands; additional AI Studio credits become easier to purchase directly.
Miro, Aug. 2026	Previously separate capabilities enter Enterprise Accelerate at unchanged existing contract prices.
Figma, Mar./Aug. 2026	Paid pooled credits and pay-as-you-go supplement seat allowances; purchased quantities later increase.
monday.com, May 2026 cohort	Defined new work-platform buyers purchase seats and an AI-credit bucket together.
Salesforce, May 2025	Agentforce adds action-based Flex Credits alongside conversation pricing.
GitHub, June 2026	Copilot premium-request allowances move toward token-based AI Credits, with migration exceptions.
Cursor, May/June 2026	Bugbot removes a seat fee and shifts to usage billing as customers reach eligible renewals.
Replit, June/July 2025	Fixed Agent checkpoint charges become effort-sensitive pricing.
Zapier, Oct. 2026	Standalone Agents subscriptions and activity metering move into common platform task billing.
HubSpot, Apr. 2026	Customer Agent charges on resolved conversations; Prospecting Agent on recommended leads.
Intercom, Mar. 2026	Fin expands the priced event from autonomous resolutions to configured successful outcomes.
Zendesk, May 2026	Automated-resolution platform introduces tiers and a funded resolution allowance.

Dates identify the transition analyzed, not necessarily a universal migration date. Sources and cohort conditions are given in the individual cases. Atlassian's separately announced December 2026 usage billing is discussed as future at the cutoff.

3 Bundling, segmentation, and the price of access

Microsoft: segment the AI offer rather than assume one enterprise price

Microsoft announced Copilot Business for qualifying SMB subscriptions in November 2025, available December 1 at $21 per user per month. Its partner announcement compared this with the existing $30 Copilot SKU and described the same productivity capabilities, eligibility tied to Business plans, and a 300-user cap.[1, 2] This was a segment-specific alternative, not a universal reduction for all customers. The research implication is to test whether different buying constraints justify a different offer, commitment, and acquisition path before assuming one AI willingness-to-pay curve spans SMB and enterprise.

Google Workspace: the same bundle can raise or lower a bill

Google's January 2025 announcement put AI into Business and Enterprise subscriptions. Its Business Standard illustration moved a prior $32 total including Gemini to $14, while the prior non-AI Standard base was $12.[3, 4] Existing AI buyers and non-AI buyers therefore faced different comparisons. The case makes cohort analysis essential: research should identify who previously purchased AI, who values the included capability, and who experiences the revision as a core-price increase rather than a saving.

Adobe: AI capabilities become package fences

Adobe's 2025 redesign renamed All Apps as Creative Cloud Pro and offered Standard with more limited AI. The US individual annual-contract, monthly-billed rates were $69.99 for Pro and $54.99 for Standard; Adobe had previously established All Apps at $59.99.[5, 6, 7] The comparison concerns a richer plan and an alternative tier, not two identical offers. The research question is which creative workflows justify premium AI access and which buyers prefer the desktop-app proposition with fewer AI entitlements.

Notion: replace an optional line item with an integrated workflow

Notion's May 2025 AI-for-work announcement moved the then-current AI toolkit into Business and Enterprise. The prior AI add-on had been separately priced by member; the revised Business offer started at $20 per member per month.[8, 9] That change encourages evaluation of the whole workflow and tier rather than a standalone AI attachment. It does not prove every buyer saves money or that every later agent capability is unmetered. Research must test package choice and migration against the specific incumbent arrangement.

Atlassian: inclusion and intensive-use monetization can coexist

Atlassian's April 2025 Rovo announcement brought search, chat, and agents into eligible paid Jira, Confluence, and Jira Service Management subscriptions through a phased rollout. Previously billed customers stopped separate Rovo invoices after onboarding.[10, 11] A September 2026 announcement described an additional usage layer, with billing scheduled for December 3, after this paper's cutoff.[12] The case supports a layered question: which AI belongs in ordinary access and which intensive work

buyers will accept as incremental consumption. Announced future charges are not evidence of current realized revenue.

Airtable: remove the seat add-on without making compute unlimited

Airtable's June 24, 2025 AI-native relaunch changed the commercial structure from an AI seat add-on to included AI capabilities and monthly credits, with additional packs available. Its billing documentation confirms the end of per-seat AI charges for the affected self-serve plans and prorated treatment of unused add-on time.[13, 14] The design separates access from capacity. The research question is how much allowance supports useful adoption and which high-use customers will buy more, rather than whether all AI should be described as free.

Grammarly: greater AI entitlement at an unchanged existing bill

Grammarly's official migration notice dates the gradual move of eligible web-purchased Premium subscriptions into Pro from March 2025. The migration preserved their bill, doubled monthly generative-AI prompts from 1,000 to 2,000, and broadened team capabilities; mobile-store subscriptions followed a different arrangement.[15] This is a packaging and entitlement change, not evidence of a price increase. The commercial hypothesis worth researching is whether a richer bundle improves adoption, collaboration, or retention enough to justify its delivery costs.

Box: remove selected caps to strengthen a premium plan

Box's March 2024 Enterprise Plus AI launch included a per-user question allocation and a company pool. In June, Box made end-user questions in Documents and Notes unlimited for Enterprise Plus, with Hubs following at general availability.[16, 17] The scope was specified interactive capabilities, not every AI workload or every plan. The case challenges the assumption that all AI needs an explicit usage charge: research may instead support inclusion as a reason to choose or retain a valuable premium subscription.

Asana: included access plus a clearer route to additional capacity

Asana widened AI Studio Basic access across eligible paid tiers in July 2025 and announced direct purchase of AI Studio Plus credit subscriptions in August. Previously, access and additional purchase had more restrictive administration or sales-assisted paths.[18, 19] This is partly a packaging and purchasing-friction change rather than a general seat-price rise. Research should distinguish whether buyers decline because of weak value, unclear allowance, or the work required to procure more capacity. Those problems call for different remedies.

Miro: fold AI capabilities into existing enterprise contracts

Miro's August 17, 2026 announcement folded formerly separate capabilities, including AI Workflows and Prototypes, into Enterprise Accelerate. Existing enterprise pricing, contract terms, and renewal dates remained unchanged.[20] This illustrates monetization through a broader license proposition rather than an immediate AI surcharge. It does not establish the effect on renewal or expansion. Buyer research should test whether the combined workflow creates value across a team and whether separate add-ons had been limiting otherwise attractive adoption.

4 Allowances, credits, and the consumption layer

Figma: add paid capacity, then change its effective quantity

Figma announced additional AI-credit purchase options in December 2025, with subscriptions starting in March 2026 alongside included seat allowances. Its subsequent FAQ records purchased credit quantities increasing at unchanged prices on August 25.[21, 22] The architecture combines subscription access and incremental capacity, while the effective credit quantity can change independently of the headline price. Research should test both ordinary use and intensive workflows. A credit package is not a fixed number of completed designs because consumption depends on the task and model.

monday.com: seats and AI capacity become a joint purchase

monday.com's official work-platform documentation identifies a new-customer cohort from May 6, 2026 that purchases seats and an AI-credit bucket together. The stated minimums are 1,000 monthly credits for Basic, 2,000 for Standard, and 3,000 for Pro. Earlier AI Blocks materials had described included credits with optional additional capacity.[23, 24] The new structure separates people from AI use while requiring both in the purchase. Research should establish whether the minimum feels like useful included capacity or compulsory expenditure on a workflow the buyer has not adopted.

Salesforce: add a finer activity unit and budget flexibility

Salesforce introduced Agentforce Flex Credits in May 2025 alongside its existing $2 conversation option. The announcement priced 100,000 credits at $500 and a standard 20-credit action at $0.10, while a Flex Agreement could reallocate commitments between user licenses and credits.[25] A conversation and an action are different units, so the figures do not establish a 95 percent like-for-like discount. The research question is which unit buyers can connect to a valuable workflow and how many chargeable actions that workflow requires.

GitHub: replace coarse requests with model-sensitive consumption

GitHub's April 27, 2026 Copilot announcement scheduled token-based AI Credits for June 1, replacing premium-request metering for affected plans while retaining exceptions for existing annual subscriptions. Base subscription rates were unchanged in the announcement.[26] The change responds to different compute intensities across coding interactions and longer autonomous work. It also changes the buyer's forecasting problem: request counts no longer suffice to estimate the bill. Research should test whether the resulting allowance, budget controls, and explanation are understandable to both developers and budget owners.

Cursor: the AI reviewer shifts from a seat to work consumed

Cursor's May 11, 2026 Bugbot announcement removed a $40 per-seat monthly subscription in favor of usage billing. Existing accounts moved at the next renewal after June 8, so an annual subscriber could remain on the prior structure into 2027.[27] This is a specific AI-code-review transition, not a claim about every Cursor plan. It changes which customers subsidize which workloads. Research should compare light and heavy review users and test whether configurable effort and spend controls make variable charging acceptable.

Replit: the price of a checkpoint ceases to be fixed

Replit's June 2025 effort-based pricing announcement replaced a fixed $0.25 Agent checkpoint charge with a cost related to the request's time and computation. More complex tasks could cost more and simple changes less; existing-user rollout began in July.[28] The case is an earlier origin of the compute-sensitive approach visible in 2026. A familiar checkpoint label does not guarantee a familiar bill. Research should investigate whether users can anticipate effort and whether the product provides enough information to make the spending decision.

Zapier: consolidate the AI bill into the platform meter

Zapier's October 2026 transition brought standalone Agents into AI by Zapier. Migrated workloads used the main task pool rather than a separate Agents subscription and activity meter; the migration guide did not set a forced shutdown date.[29, 30] Consolidation can simplify purchasing, but it does not guarantee lower workload cost because task consumption still depends on model and tool activity. Research should test the complete workflow bill and the value of a common budget, rather than assume that removing an add-on automatically improves economics for every account.

5 Outcome pricing is a definition, not just a label

HubSpot: change the rate and narrow the paid event

Effective April 14, 2026, HubSpot moved Breeze Customer Agent from $1 per conversation to $0.50 per resolved conversation. Prospecting Agent moved from recurring charges for enrolled contacts to $1 per lead recommended for outreach, with credits used as the payment mechanism.[31] These are changes in what earns a charge, not simply lower unit prices. Research should establish whether the defined completion corresponds to what the buyer values. A recommended lead is not a guaranteed sale, just as a conversation is not necessarily a resolved problem.

Intercom: broader successful work can widen the paid event

Intercom's March 12, 2026 Fin announcement expanded the metric from fully autonomous resolutions to configured successful outcomes, including Procedures designed to hand off to a human. Its outcome documentation identifies the standard $0.99 charge and relevant exclusions.[32, 33] This can broaden billable work without a higher nominal rate. Research should test whether buyers recognize the partial workflow as valuable and understand the distinction between a successful configured handoff and an unsuccessful escalation. "Outcome-based" does not guarantee identical success definitions across suppliers.

Zendesk: classify resolutions and fund an allowance

Zendesk's official documentation identifies resolution tiers introduced May 18, 2026, with an allowance from which the cost is drawn. It distinguishes assisted escalations and unverified containment from verified resolution, while customers on the prior platform remain there until moved.[34, 35] Public tier-specific numerical rates were not established for this comparison. The analytical point is the classification and allowance architecture, not a fabricated before-and-after percentage. Research should determine whether customers accept the verification rule and can connect the priced result to their own operational value.

6 Six findings that matter for a pricing decision

AI monetization is layered rather than uniformly seatless

The sample does not support a universal migration away from subscriptions. Several businesses broadened what a paid seat or tier includes. Others retained a subscription and added consumption, changed an add-on, or priced a completed event. The commercially useful question is which layers correspond to access, differentiated workflow value, intensive use, and assurance in the buyer's situation.

For a portfolio company, a hybrid can be appropriate when baseline access is broadly useful but cost or value varies sharply with activity. Inclusion can be appropriate when it improves the core offer enough to justify its cost. Neither is established by observing a competitor's decision. Research must connect the architecture to the target segment and business objective.

The denominator can move more than the headline rate

Conversation, resolution, action, handoff, request, token, and credit name different economic objects. A headline comparison is meaningful only when scope, quantity, billing period, success rule, and commitment are comparable. A nominal reduction can accompany a narrower event; an unchanged rate can apply to a broader event.

Use a bill function rather than a single rate in competitive analysis:

$$B = F + \sum_{k} r_k q_k,$$

where F includes fixed commitments, r_k is the charge for a defined event, and q_k is the quantity under the offer. The quantities may change when the tariff changes. Before copying a price, research the events the customer understands and the behavior the charge will induce.

An included allowance is an allocation decision

An allowance determines which use is funded by the subscription, who subsidizes intensive consumption, and when a buyer encounters another purchase. Its economic effect depends on the distribution of usage and costs, not only the mean. Pooling can make spare capacity useful across a team while also changing who controls spend.

Research low-use, typical, and intensive accounts. Ask whether the allowance permits a meaningful workflow, whether the buyer can forecast depletion, and what happens when the limit is reached. A usage dashboard does not by itself establish that the customer can understand or budget the proposed charge.

Packaging changes the relevant reference price

The prior buyer arrangement matters. A customer paying separately for AI may see integration as a saving. A customer who never chose the add-on may see a higher core bill. A customer losing an entitlement may evaluate premium as a repurchase rather than new value. These differences can be hidden in average contract-value reporting.

Segment research by incumbent offer and buying occasion as well as by firm size or industry. A new prospect chooses among current alternatives; an existing customer also compares the proposal with prior

access, promises, and price. The migration case can therefore require a different offer or timetable from new acquisition.

Outcome pricing moves risk and requires a verification contract

If unsuccessful attempts consume cost but produce no charge, the supplier bears performance risk. If the definition includes work the customer does not recognize as success, the buyer can perceive a different risk. The commercial design must specify the event, acceptance, time window, exclusions, and correction policy.

Research that definition before asking about the price. Give buyer and supplier teams the same sample work and ask each to calculate the bill. Then test whether the buyer considers the counted result useful. Agreement about arithmetic is necessary for a reproducible charge; it is not sufficient evidence of willingness to pay.

Public prices do not reveal the demand curve or the earnings result

These sources establish commercial decisions and, in some cases, the supplier's stated rationale. They generally do not isolate price-induced conversion, retention, quantity, concessions, or margin. An adoption statistic reported alongside an announcement is not evidence that the new price caused adoption.

For an investor, the gap matters. A competitor's published offer can frame a hypothesis, but it cannot establish the pricing upside in the portfolio company. That requires the company's own buyer evidence, contract and transaction data, delivery-cost distribution, and a credible implementation and measurement plan.

7 The research brief to commission before copying a model

The comparative evidence suggests five linked workstreams. The deliverable should be a commercial decision with a financial case, rather than a catalogue of competitors.

1. **Buyer and workflow research.** Identify the economic buyer, purchase trigger, current alternative, expected outcome, and conditions needed for adoption. Separate user enthusiasm from budget approval.
2. **Offer-choice research.** Test complete packages, units, rates, allowances, and commitments against realistic alternatives, including no purchase or the current process. Measure comprehension before interpreting preference.
3. **Usage and cost analysis.** Join attempted work, accepted output, model and tool consumption, review, infrastructure, and support. Estimate the economic distribution across accounts and plausible high-use conditions.
4. **Financial and migration modeling.** Forecast net realization, substitution, contraction, retention, eligibility, and timing. Keep annualized opportunity separate from current-year earnings.
5. **Commercial validation.** Put a defined offer into a bounded paid test where feasible. Observe purchase, use, contribution, and continued demand, retaining the exact terms and limitations of the pilot.

Prioritize the uncertainties that can change the recommendation. If buyers cannot explain the credit bill, more precise cost estimates will not resolve the purchase barrier. If the offer is understood but high-use accounts lose money, another awareness survey will not fix the cost distribution. Match the research to the unresolved decision.

8 Conclusion: research your buyers before borrowing a pricing model

The 20 cases show a market experimenting with the object of purchase. Some suppliers include more AI in subscriptions. Some differentiate packages or segment access. Others meter effort, consolidate platform consumption, or redefine a successful outcome. The evidence supports several architectures, not one universal AI tariff.

The unanswered questions belong to the individual business: which workflow creates value, which buyer funds it, which unit is understandable, which package buyers choose, what expenditure they will accept, and whether the resulting usage supports contribution. Public competitor data cannot supply those observations. Primary pricing research and operating data are needed to turn a market comparison into an evidence-based recommendation.

Curly's Consulting's **Pricing Transformation** combines buyer research and willingness-to-pay analysis, monetization design, size-of-prize modeling, and direct execution support. Commission a scope around the commercial choice the management team needs to make. If materiality is still uncertain, the **Pricing Upside Report** provides conditional revenue and EBITDA scenarios and priorities for the next step; original buyer research and implementation are part of the Transformation scope.

Turn the market evidence into an answer for your business

Request a Pricing Transformation quote for original buyer evidence and a commercial recommendation. To assess the opportunity first, buy the Pricing Upside Report.

References

[1] Microsoft. *Microsoft Ignite 2025: Copilot and agents built to power the Frontier Firm*. November 18, 2025. https://www.microsoft.com/en-us/microsoft-365/blog/2025/11/18/microsoft-ignite-2025-copilot-and-agents-built-to-power-the-frontier-firm/.

[2] Microsoft. *November 2025 Partner Center announcements: Microsoft 365 Copilot Business*. November 18, 2025. https://learn.microsoft.com/partner-center/announcements/2025-november.

[3] Google Workspace. *Google Workspace enables the future of AI-powered work for every business*. 2025-01-15; timing updated 2025-04-28. https://workspace.google.com/blog/product-announcements/empowering-businesses-with-AI.

[4] Google Workspace. *Expanding Google AI to more of Google Workspace*. 2025-01-15. https://workspaceupdates.googleblog.com/2025/01/expanding-google-ai-to-more-of-google-workspace.html.

[5] Adobe. *Meet Creative Cloud Pro: new tools and expansive creative controls*. 2025-05-15. https://blog.adobe.com/en/publish/2025/05/15/meet-creative-cloud-pro-new-tools-expansive-creative-controls.

[6] Adobe. *Changes to Creative Cloud plans for individuals*. Updated 2026-02-12. https://helpx.adobe.com/account/individual/subscriptions-and-plans/plan-types-and-eligibility/changes-to-individual-plan.html.

[7] Adobe. *AI, Creative Cloud, and our new pricing update*. 2023-09-13. https://blog.adobe.com/en/publish/2023/09/13/ai-creative-cloud-release-pricing-update.

[8] Notion. *Notion AI for work: official product and pricing announcement*. 2025-05-13. https://www.notion.com/blog/notion-ai-for-work.

[9] Notion. *Introducing Q&A*. 2023-11-14. https://www.notion.com/blog/introducing-q-and-a.

[10] Atlassian. *Rovo for all: Team 25 announcement*. 2025-04-09. https://www.atlassian.com/blog/announcements/team25-rovo-for-all.

[11] Atlassian. *Understand Rovo billing and managing costs in Atlassian Cloud*. Updated 2025-09-26. https://support.atlassian.com/rovo/kb/understand-rovo-billing-and-managing-costs-in-atlassian-cloud/.

[12] Atlassian. *Expanded usage-based pricing*. 2026-09-01. https://www.atlassian.com/blog/company-news/expanded-usage-based-pricing.

[13] Airtable. *The AI-Native Airtable Has Arrived: A Letter from our CEO*. 2025-06-24. https://www.airtable.com/newsroom/introducing-the-ai-native-airtable.

[14] Airtable. *Airtable AI billing*. Search result says updated October 2, 2026; opened page says one month ago. Explicit historical transition date June 24, 2025.. https://support.airtable.com/articles/3378106230-airtable-ai-billing.

[15] Grammarly / Superhuman. *Updating the Grammarly Premium plan to the Grammarly Pro plan*. No publication date shown; text explicitly dates migration beginning March 2025.. https://support.grammarly.com/hc/en-us/articles/33347704601613-Updating-the-Grammarly-Premium-plan-to-the-Grammarly-Pro-plan.

[16] Box. *Box Expands its Collaboration with Microsoft with New Azure OpenAI Service Integration*. 2024-03-05. https://www.boxinvestorrelations.com/news-and-media/news/press-release-details/2024/Box-Expands-its-Collaboration-with-Microsoft-with-New-Azure-OpenAI-Service-Integration/default.aspx.

[17] Box. *Box Announces Powerful Enhancements to Box AI for Intelligent Content Management*. 2024-06-27. https://www.boxinvestorrelations.com/news-and-media/news/press-release-details/2024/Box-Announces-Powerful-Enhancements-to-Box-AI-for-Intelligent-Content-Management/default.aspx.

[18] Asana. *AI Studio Basic is coming soon for all users on eligible paid plans*. 2025-07-01. https://forum.asana.com/t/ai-studio-basic-is-coming-soon-for-all-users-on-eligible-paid-plans-no-admin-setup-needed/1076628.

[19] Asana. *Introducing AI Studio Plus for direct purchase*. 2025-08-11. https://forum.asana.com/t/introducing-ai-studio-plus-for-direct-purchase/1085642.

[20] Miro. *New Enterprise pricing: Miro's adding more AI capabilities to Enterprise licenses at no additional cost*. 2026-08-17. https://miro.com/blog/enterprise-pricing-miro/.

[21] Figma. *Updates to AI credits in Figma*. 2025-12-09. https://www.figma.com/blog/updates-to-ai-credits-in-figma/.

[22] Figma. *AI credit updates FAQ*. Undated FAQ; explicit changes Aug 25 and Oct 6, 2026. https://help.figma.com/hc/en-us/articles/42614902212887-AI-credit-updates-FAQ.

[23] monday.com. *The pricing model for monday AI portfolio*. No original publication date verified. Explicit applicability/effective boundary May 6, 2026; retrieved before cutoff.. https://support.monday.com/hc/en-us/articles/35277848309394-The-pricing-model-for-monday-AI-portfolio.

[24] monday.com. *monday.com Announces AI Vision to Empower Businesses to Scale*. 2025-02-10. https://ir.monday.com/news-and-events/news-releases/news-details/2025/monday.com-Announces-AI-Vision-to-Empower-Businesses-to-Scale/default.aspx.

[25] Salesforce. *Salesforce Introduces New Flexible Agentforce Pricing to Accelerate the Digital Labor Revolution*. 2025-05-15. https://www.salesforce.com/news/press-releases/2025/05/15/agentforce-flexible-pricing-news/.

[26] GitHub. *GitHub Copilot is moving to usage-based billing*. 2026-04-27. https://github.blog/news-insights/github-copilot-is-moving-to-usage-based-billing/.

[27] Cursor. *Updates to Bugbot for Teams and Individuals*. 2026-05-11. https://prod.cursor.com/blog/may-2026-bugbot-changes.

[28] Replit. *Introducing Effort-Based Pricing for Replit Agent*. 2025-06-18. https://replit.com/blog/effort-based-pricing.

[29] Zapier. *Zapier Agents is now AI by Zapier*. 2026-10-06. https://zapier.com/blog/zapier-agents-is-now-ai-by-zapier/.

[30] Zapier. *Migrating from Agents to AI by Zapier*. Official documentation; dated effective change discussed in text. https://help.zapier.com/hc/en-us/articles/47402591569805-Migrating-from-Agents-to-AI-by-Zapier.

[31] HubSpot. *HubSpot's Customer Agent and Prospecting Agent: Now you pay when the task is complete*. 2026-04-13. https://www.hubspot.com/company-news/hubspots-customer-agent-and-prospecting-agent-now-you-pay-when-the-task-is-complete.

[32] Intercom. *From resolutions to outcomes: Evolving how Fin delivers value*. 2026-03-12. https://www.intercom.com/blog/from-resolutions-to-outcomes-evolving-how-fin-delivers-value.

[33] Intercom. *Fin AI Agent outcomes*. 2026-09-23. https://www.intercom.com/help/en/articles/8205718-fin-ai-agent-outcomes.

[34] Zendesk. *About automated resolution tiers*. Documentation edited Sep 30, 2026; explicit introduction May 18, 2026. https://support.zendesk.com/hc/en-us/articles/9570369117338-About-automated-resolution-tiers.

[35] Zendesk. *About the automated resolutions platform (prior to resolution tiers)*. Explicitly prior to May 18, 2026. https://support.zendesk.com/hc/en-us/articles/5352026794010-About-the-automated-resolutions-platform-prior-to-resolution-tiers?region=united-states.

About the author. Ammanuel "Tom" Selameab, CPP, is a Certified Pricing Professional and founder of Curly's Consulting. He helps private equity investors and portfolio-company leaders answer pricing and monetization questions through primary buyer research, commercial analysis, and financial modeling.

Curly's products | tom@curlysconsulting.com

© 2026 Curly's Consulting. All rights reserved.