

FRONTIER.

FRONTIER / NEW ANALYSIS



Salesforce and monday.com Put AI Agents on Usage Meters, Breaking SaaS's Seat-to-Value Link

Salesforce is selling Agentforce through credits, conversations, users and resolutions, while monday.com has introduced seats plus AI credits—a broader shift from pricing human access to metering machine-executed work.

NEW ANALYSIS

Built to change a decision.

Salesforce is selling Agentforce through credits, conversations, users and resolutions, while monday.com has introduced seats plus AI credits—a broader shift from pricing human access to metering machine-executed work.

EDITION	New Analysis
AUTHOR	Frontier Research
PUBLISHED	October 19, 2026

This PDF is a publication rendering of the canonical Frontier News article. The web article remains the live source of truth; analysis, figures, and source citations travel with this edition.

What changed

Enterprise AI pricing is moving beyond the seat, but not simply replacing it. Salesforce currently offers several Agentforce buying models at once: Flex Credit consumption pricing, per-conversation pricing, per-user licenses and resolution-based pricing.[1] Its rate card lists 100,000 Flex Credits for \$500, with a standard action consuming 20 credits and a voice action consuming 30. At the listed rates, those units calculate to \$0.10 and \$0.15 respectively, although the cost of an end-to-end business process will depend on how many billable actions it requires.[1]

monday.com made the hybrid structure explicit in its May 11, 2026 first-quarter disclosure, announcing a seats-plus-credits model for its AI portfolio.[2] The company’s support documentation says customers covered by that model buy seats for people and credits for AI usage.[3] Management characterized the change as a move toward consumption-based pricing and argued that its business can grow as AI performs more work for customers; that is a vendor claim about the model’s economics, not an independently established outcome.[2]

Meta has disclosed another variation. The company told investors that Meta Business Agent monetization would combine subscriptions with volume-based token pricing.[4] Independently, TechCrunch reported that Respond.io prices according to conversation volume rather than seats, limiting the direct revenue effect when AI handles work previously performed by people.[5] The Information has likewise reported that enterprise software providers are experimenting with consumption- and outcome-based structures because AI agents weaken the relationship between software use and employee headcount.[6]

The event, therefore, is not the universal death of per-seat software. It is the emergence of multiple commercial meters alongside seats. Salesforce still sells a per-user Agentforce option,[1] and Microsoft told Fortune in March 2026 that its customers currently preferred per-user pricing.[7] Forkast's October 16, 2026 analysis described four competing approaches to agent monetization, but its categorical framing that per-seat pricing is dying goes beyond the evidence.[8] The verified pattern is a transition toward hybrid contracts that can charge separately for human access and machine-executed work.

Why it matters

For CIOs, the central buying question is changing from “How many employees need access?” to “What work will the system perform, which events create charges, and how variable will the resulting bill be?” Seat counts are relatively observable before deployment. Agent consumption depends on workload volume, process design and the vendor's definition of a billable unit. A credit, token, action, conversation and completed resolution can each represent a different scope of work, so their published prices are not directly comparable.[1][4]

That distinction invalidates simplistic price comparisons. Forkast cited a claimed 40-to-50-times price advantage for Meta, but the underlying units do not support a reliable like-for-like conclusion: Meta described subscription and token-volume pricing, while Salesforce lists actions, conversations, users and resolutions.[8][1][4] A lower price per token cannot establish a lower cost per completed workflow unless the buyer also knows how many tokens, actions, retries and human interventions the workflow requires. The analytical implication is that CIOs must compare vendors at the level of a defined workload or accepted outcome, not at the level of the vendor's preferred meter.

The budget model changes with it. Under a seat contract, automation that reduces the number of human users can reduce the software bill at renewal. Under consumption pricing, the same automation may increase usage because the agent performs more transactions. monday.com explicitly presents that alignment as a benefit to its business, saying it can grow as AI does more customer work.[2] Buyers should read that statement as a clear description of the supplier incentive: greater machine activity can generate greater vendor revenue even when employee headcount does not rise.

Gartner estimated that up to \$234 billion—about 20% of projected enterprise application SaaS spending—could be exposed to “agentic arbitrage” by 2030 and said the development breaks the traditional relationship between seats and licenses.[9] Fortune separately reported that CIOs want suppliers to move away from per-seat fees and that agentic AI is pressuring per-user pricing.[10] Those findings support the direction of travel, but Microsoft’s reported customer preference for per-user pricing shows why CIOs should not assume one successor model will dominate immediately.[7]

The decision model must consequently add workload forecasts, unit-definition scrutiny and consumption-risk controls to conventional license management. CIOs need to determine whether a contract rewards useful automation or merely monetizes activity. They also need to test whether an “outcome” is objectively accepted by the enterprise, whether retries are billable, and whether a vendor can change the work represented by a credit. These are analytical requirements created by the documented pricing structures; they are not claims that every vendor currently applies the same terms.

Frontier take

The decisive commercial battleground for enterprise agents will be control of the billing unit, not the nominal price attached to it. That is the strongest conclusion supported by the coexistence of Salesforce’s credits, actions, conversations, users and resolutions; monday.com’s seats-plus-credits model; and Meta’s proposed mix of subscriptions and token volume.[1][3][4]

A seat has an imperfect but familiar denominator: a person authorized to use the product. Agent pricing fractures that denominator into vendor-specific abstractions. Suppliers can present each abstraction as closer to value, but a billing unit is not an outcome merely because it is generated by AI. An action may be one step in a process. A conversation may end without resolution. Tokens measure processing volume rather than business completion. Even resolution pricing requires agreement on what counts as resolved. The unit’s definition therefore determines whether automation transfers economic value to the buyer or converts that value into an expanding variable charge.

Per-user licensing will persist where buyers value predictability or where human access remains a reasonable usage proxy. Salesforce's active per-user option and Microsoft's reported customer feedback are direct counterevidence to claims of an immediate, universal collapse.[1][7] The more credible structural shift is toward layered contracts: seats for people, consumption meters for agents and, in some cases, charges tied to conversations or outcomes. Independent reporting already documents experimentation with those alternatives across the market.[6][10]

CIOs should therefore treat pricing architecture as part of system architecture. A workflow design that uses more agent actions, longer conversations or greater token volume can alter economics even when the software's sticker rate does not change. Procurement cannot evaluate these contracts separately from process engineering, and engineering cannot optimize an agent without visibility into the commercial meter. The winning enterprise discipline will be the ability to translate every proposed billing unit into cost per accepted unit of work—and to reject contracts where that translation cannot be audited.

Three moves for CIOs

1. — Build a workload bill of materials before comparing agent prices Select representative workflows and map every expected billable event: human seats, agent actions, conversations, tokens, retries and accepted outcomes. Require each shortlisted vendor to price the same workload under every available model, then calculate expected cost and a high-volume scenario per enterprise-accepted outcome. Keep vendor-native units visible rather than converting them through unsupported equivalencies.
 - Decision trigger: Make this analysis mandatory whenever a proposal includes more than one billing unit, combines seats with variable consumption, or cannot state the expected cost of a defined end-to-end workflow.
 - Why now: Salesforce already exposes several simultaneous meters, while monday.com and Meta have documented different hybrid structures.[1][3][4] Comparing their headline unit prices would confuse unlike scopes of work. A workload bill of materials gives the CIO a common denominator without pretending that an action, token, conversation and resolution are equivalent.
2. — Contract for semantic stability, not only price protection Define each billable unit in the order form, including treatment of retries, failed actions, test activity, duplicate events and human escalation. Require auditable usage records, advance notice of measurement changes and a repricing or termination right if the vendor changes how credits are consumed. Add budget alerts and a hard approval gate before overages can exceed the enterprise's agreed ceiling.
 - Decision trigger: Escalate the contract when the supplier can change credit consumption, redefine a resolution, charge for unsuccessful work or permit variable spend above the approved ceiling without affirmative customer approval.

- Why now: The documented models meter fundamentally different events.[1][4] As machine activity becomes a revenue driver, nominal rate protection is insufficient if the supplier can alter how much work one credit represents. Locking down the unit's meaning protects the economics of automation rather than merely preserving a published price.
- 3. — Run seat and consumption models in parallel before committing During a controlled production period, calculate the same workload under per-user, credit, conversation and outcome-based alternatives wherever the vendor offers them. Compare not only average cost but also cost variability and the share of charges associated with unsuccessful or unaccepted work. Favor the structure whose bill most closely tracks an outcome the enterprise can verify.
- Decision trigger: Require a parallel-model test before renewal when agent adoption could reduce human seats, when usage volume is materially uncertain, or when more than 10% of measured charges arise from retries, failures or work the enterprise does not accept as complete.
- Why now: Per-seat pricing has not disappeared: Salesforce retains a per-user option, and Microsoft reported continued customer preference for that model.[1][7] CIOs can use the coexistence of old and new structures as negotiating leverage. Testing both exposes whether a usage model rewards productive automation or simply shifts demand uncertainty from the vendor to the customer.

Sources

- Forkast — Agent Monetization Shift: Per-Seat Pricing Is Dying [8]
- Gartner — \$234 Billion in Enterprise Application Software Spend at Risk [9]
- Salesforce — Agentforce Pricing [1]
- monday.com — First Quarter 2026 Results [2]
- monday.com — AI Portfolio Pricing Model [3]
- Meta — Q2 2026 Follow-Up Call Transcript [4]
- The Information — Can AI Tollgates Become the New SaaS Model? [6]
- Fortune — CIOs and CTOs Take a Harder Line With SaaS Vendors [10]
- TechCrunch — Respond.io Uses Conversation-Volume Pricing [5]
- Fortune — Microsoft Says Customers Currently Want Per-User Pricing [7]

SOURCES / CITATIONS

Trace the argument.

Every source below is linked. Source numbering follows first appearance in the report.

- 01** **Salesforce — Agentforce Pricing**
<https://www.salesforce.com/agentforce/pricing/>
- 02** **monday.com — First Quarter 2026 Results**
<https://monday.com/p/press-release/monday-com-announces-first-quarter-2026-results/>
- 03** **monday.com — AI Portfolio Pricing Model**
<https://support.monday.com/hc/en-us/articles/35277848309394-The-pricing-model-for-monday-AI-portfolio>
- 04** **Meta — Q2 2026 Follow-Up Call Transcript**
https://s21.q4cdn.com/399680738/files/doc_financials/2026/q2/META-Q2-2026-Follow-Up-Call-Transcript.pdf
- 05** **TechCrunch — Respond.io Uses Conversation-Volume Pricing**
<https://techcrunch.com/2026/06/15/malaysias-respond-io-raises-62-5m-eyes-acquisitions-in-north-america-and-europe/>
- 06** **The Information — Can AI Tollgates Become the New SaaS Model?**
<https://www.theinformation.com/articles/can-ai-tollgates-become-new-saas-model>
- 07** **Fortune — Microsoft Says Customers Currently Want Per-User Pricing**
<https://www.fortune.com/2026/03/09/microsoft-copilot-cowork-ai-agents-anthropic-e7-m365-saas/>
- 08** **Forkast — Agent Monetization Shift: Per-Seat Pricing Is Dying**
<https://forkast.news/agent-monetization-shift-per-seat-pricing-is-dying/>
- 09** **Gartner — \$234 Billion in Enterprise Application Software Spend at Risk**
<https://www.gartner.com/en/newsroom/press-releases/2026-07-01-gartner-says-us-dollars-234-billion-in-enterprise-application-software-spend-is-at-risk-from-agentic-artificial-intelligence>
- 10** **Fortune — CIOs and CTOs Take a Harder Line With SaaS Vendors**
<https://www.fortune.com/2026/04/08/amid-the-saaspocalypse-cios-and-ctos-take-a-harder-line-with-their-vendors/>

FRONTIER.

TAKE THE NEXT STEP

Push this further.

Research is useful when it changes a decision. Bring us the question, the constraint, or the argument you want tested.

RESEARCH / BRIEFINGS

inquiries@thisisfrontier.ai

Briefings, advisory work, executive discussions, and research inquiries.

REPRINTS / LICENSING

reprints@thisisfrontier.ai

Reprints, enterprise distribution, excerpts, and licensing.

FRONTIER

www.thisisfrontier.ai

Research, Data Insights, News, and the live Frontier corpus.

All rights reserved. Contact Frontier for redistribution or reprint rights.

SEE FURTHER. MOVE SOONER.