

# The Metric Keeps Changing: What AI Customer Service Agents Actually Cost

Unvarnished Reviews · TCO Teardown · July 2026 · [unvarnishedreviews.com](https://unvarnishedreviews.com) · Free · No subscription

**The contradiction: Every vendor in this comparison says the same thing: "You only pay when it works." None of them agree on what "it" is, what "works" means, or how much it costs when it does. Zendesk redefined its own billing unit twice in five months. HubSpot switched pricing models entirely with twelve days' notice and no grandfathering. Salesforce is running three incompatible pricing models simultaneously because it couldn't get customers to accept any single one. One industry researcher found 75% of vendors selling AI agents have no systematic pricing approach at all. This is not a mature market pricing a mature product. This is four vendors improvising in public, and the buyer is the one who has to reconcile the bill.**

## The verdict up front

Every platform here markets "outcome-based pricing" as the industry's answer to AI cost anxiety: Pay for results, not access. The pitch is genuinely appealing. The execution is the problem — because "resolution," "conversation," and "credit-consuming action" are four different units of work priced by four different vendors, none of them comparable to the others, and at least two of them have already changed definition mid-flight. Before you can compare cost, you have to reverse-engineer what you're actually being charged for. That reverse-engineering is this report.

## HubSpot Breeze: The pricing model that changed with twelve days' notice

On April 2, 2026, HubSpot emailed customers a subject line that undersold what was coming: "Changes coming to Breeze prospecting agent on April 14." Twelve days later, both AI agents — Customer Agent and Prospecting Agent — moved to a new pricing model, for every customer, new and existing, with no grandfathering.

The new rates: Customer Agent dropped from \$1.00 per conversation to \$0.50 per **resolved** conversation. Prospecting Agent moved from a flat monthly charge per enrolled contact to \$1.00 per qualified lead recommended for outreach. HubSpot's Chief Customer Officer, Jon Dick, framed the change plainly: "You pay when it works, full stop." The company's own announcement put it as a principle: "AI should be priced on the value it delivers, not the compute it consumes."

The included credit allowances stayed identical across tiers — 500 credits on Starter, 3,000 on Professional, 5,000 on Enterprise — which means the same allowance now covers roughly twice as many outcomes as it did under the old per-conversation model, since a resolution costs half what a conversation used to. On paper, most customers' bills fall.

The catch is in the word "resolved." Per independent analysis of the rollout, the headline 50% cut is real for straightforward cases, but HubSpot's own definition of what counts as "resolved" narrows the charge base in ways that don't universally favor the customer: Long-running email threads and multi-issue conversations can,

in specific scenarios, quietly cost more under the new model than the old one. The rate is 50 credits per resolution, which at HubSpot's \$9.00 per 1,000 credits works out to roughly \$0.45–\$0.50 per resolution depending on overage tier.

### **Intercom Fin: The one that got there first, and stayed simple**

Intercom's Fin has run pure per-resolution pricing since before this was a trend, and reporting puts it past \$100 million in annual recurring revenue on that model as of this year — evidence that outcome-based pricing can work at scale when a vendor commits to one model instead of hedging across several. The rate, per independent pricing analysis, is \$0.99 per resolution. Intercom's structure is a hybrid: The core helpdesk seats are priced per agent per month in the traditional way, and Fin's AI resolutions are billed separately, on top, at the per-resolution rate. Of the four vendors here, Intercom's model is the least confusing to explain — which is itself notable, given the field.

### **Zendesk: The billing unit that was redefined twice in five months**

Zendesk's pricing model changed structurally in January 2026 and again in May 2026, and the two changes point in different directions for the buyer.

In January 2026, Zendesk began auto-billing resolution overages above committed volume — uncapped, with no grace period and no prior-month warning, which independent teardowns describe as the most-cited billing surprise in 2026 customer reviews. Then, on May 18, 2026, Zendesk restructured "automated resolutions" into three tiers, per its own help documentation: **Assisted Escalations** and **Contained Resolutions** are now free and don't draw from any allowance; only **Verified Resolutions** — cases fully closed by AI with no human escalation — count against the billed allotment. Zendesk frames this as tiered pricing that "reflects the value the AI agent provided." A less generous read: The billing unit itself moved twice in less than half a year, and Zendesk does not publish the resulting overage rate. Third-party teardowns report figures ranging from \$1.20–\$1.50 per Verified Resolution on committed volume up to \$2.00 pay-as-you-go — a real spread, and one buyers should get in writing rather than assume.

Two more charges stack on top: A \$50-per-agent-per-month Copilot add-on (the assistive, human-facing layer, separate from the autonomous resolution billing), and the base Suite subscription itself, list-priced \$19–\$115 per agent per month across five tiers, with the top tier now sitting behind a "Contact Sales" button rather than a published price. One independent teardown modeled a 10-agent Suite Professional team with Copilot enabled at moderate AI volume landing around \$10,650 per month — roughly 8x the base subscription. The AI layer, not the seats, is where the money goes now.

### **Salesforce Agentforce: Three pricing models running at once, because none of them worked alone**

Agentforce's pricing history is the clearest evidence in this report that outcome-based AI pricing is still being invented in real time, not executed from a stable playbook.

**Model one, launched late 2024:** A flat \$2 per conversation, regardless of how many backend actions a conversation required. It flopped commercially — of roughly 5,000 Agentforce deals in the model's first two quarters, only 3,000 converted to paying customers, per industry reporting. The stated reason: Buyers couldn't model their costs, since a single customer query could trigger eight or more backend processes under one

"conversation" charge, and the flat rate priced out smaller organizations without large AI budgets outright.

**Model two, May 2025:** Flex Credits — \$500 per 100,000 credits (roughly \$0.005 per credit), with a standard agent action consuming 20 credits (~\$0.10) and a voice action 30 credits (~\$0.15). More granular and closer to actual work performed, but still consumption-based, which independent analysis notes is still fundamentally hard to forecast for enterprise procurement — every record update, query, or retrieval step burns credits, so a single user request can consume several actions without a clean way to predict how many in advance.

**Model three, late 2025:** Per-user licensing under the Agentic Enterprise License Agreement — \$125+ per user per month, with an Agentforce 1 Edition at \$550 per user per month (annual commitment) bundling 1 million Flex Credits per year. This gave CFOs a number they could budget and CIOs a contract shape they understood.

Salesforce did not retire the earlier models when it introduced the third — **all three now run simultaneously**, though Flex Credits and per-conversation billing cannot both be active in the same org; a buyer must choose one consumption model, then can layer per-user licensing on top. The multi-model approach appears to be working better than any single model alone: Agentforce reached \$540 million in ARR by Q3 FY2026, up 330% year-over-year, with 18,500 total deals closed and 9,500 paid — though that also means roughly half of Agentforce deals still don't convert to paying usage, and only about 8% of Salesforce's 150,000+ customer base has adopted the product so far. Worth noting separately: Advanced Agentforce deployments commonly require Data Cloud (recently rebranded Data 360), which independent analysis puts at pushing real first-year costs into the \$150,000–\$600,000 range for mid-market teams — a cost most per-action or per-conversation math never surfaces.

One more naming wrinkle for anyone comparing quotes: Salesforce Sales Cloud itself has been rebranded "Agentforce Sales" in current product listings — the same kind of quiet rename that's currently confusing CNAPP buyers over Prisma Cloud (see our companion report on Wiz vs. Orca vs. Prisma Cloud). If your last quote or comparison predates the rename, confirm you're pricing the current product.

**The math: Why we are not giving you one number**

A responsible comparison would show one total per vendor for an identical workload. We can't do that in good faith here, and saying so is itself the finding. "Resolution," "conversation," and "credit-consuming action" are not interchangeable units — a HubSpot resolution and a Zendesk Verified Resolution may represent very different amounts of underlying work, and a Salesforce "conversation" under the legacy model could contain what another vendor would bill as four or five separate actions. Below is what each vendor's own reported unit price computes to under one illustrative, clearly-labeled workload: 5,000 AI-handled customer interactions per month, support-side only (Salesforce's Prospecting-side and Data Cloud costs are excluded from this comparison; they don't have an equivalent on the other three platforms).

| HubSpot Breeze Customer Agent | \$0.50 per resolved conversation | \$2,500 | \$30,000 |
|-------------------------------|----------------------------------|---------|----------|

|                                                                             |                                                |                                   |            |
|-----------------------------------------------------------------------------|------------------------------------------------|-----------------------------------|------------|
| Intercom Fin                                                                | \$0.99 per resolution                          | \$4,950                           | \$59,400   |
| Zendesk (Verified Resolutions + Copilot, 100-agent team)                    | ~\$1.20–\$2.00/resolution + \$50/agent Copilot | ~\$12,000 (resolutions + Copilot) | ~\$144,000 |
| Salesforce Agentforce, legacy per-conversation                              | \$2.00 per conversation                        | \$10,000                          | \$120,000  |
| Salesforce Agentforce, Flex Credits (illustrative, ~4 actions/conversation) | ~\$0.40 per conversation equivalent            | ~\$2,000                          | ~\$24,000  |

Read the Salesforce rows together, not separately: The same workload can price out roughly 5x apart depending purely on which of Salesforce's own three models you land on and how action-dense your actual conversations turn out to be — a spread no other vendor in this report produces internally. That unpredictability, not the sticker price of any one row, is the finding to take into a negotiation.

## The decision framework

**Choose HubSpot Breeze if:** You're already on HubSpot's platform and your support conversations are largely straightforward — the per-resolution rate is competitive and the credit allowance now stretches further. Get HubSpot's exact definition of "resolved" in writing before you model savings against the old per-conversation bill; the scenarios where it costs more are the ones your CFO will ask about later.

**Choose Intercom Fin if:** You want the simplest pricing structure to explain internally and forecast against. It's the one model in this report that hasn't changed shape, and \$100M+ ARR on a single pricing model is real evidence it works at scale.

**Choose Zendesk if:** You're already committed to the Zendesk platform and can negotiate a committed-volume rate and a spend cap into your contract. Do not accept the pay-as-you-go, uncapped default if your resolution volume is likely to grow — that's precisely the scenario the January 2026 auto-billing change was built to monetize.

**Choose Salesforce Agentforce if:** You want the per-user AELA model specifically for budget predictability, and you have Salesforce administrators on staff to manage the underlying complexity. Avoid the legacy per-conversation model and the pure Flex Credits model for anything except a bounded pilot — both are the versions Salesforce's own commercial history shows customers rejecting at scale.

**Everyone:** Before signing, ask each vendor for a real usage report from a comparable existing customer, not a projection. Every model in this report is complex enough that the vendor's own sales estimate and your actual bill can diverge significantly once real conversation complexity enters the picture.

## The bottom line

Outcome-based AI pricing is being sold as the mature, principled alternative to per-seat SaaS — pay for value, not access. The principle is sound. The industry executing it is not yet stable: Two of the four vendors here have changed their own pricing model's definition within the past six months, and the market leader is running three models at once because none of them worked alone. None of that makes outcome-based pricing wrong. It makes it a category where the contract terms matter more than the headline rate, and where "we'll get back

to you with usage-based numbers" is not an answer to accept without a comparable customer's real invoice behind it.

### **Sources and verification notes**

HubSpot Breeze pricing change (April 14, 2026 effective date, April 2 announcement, no grandfathering, credit allowances, "resolved" definition caveats): SaaStr, resolve247.ai, almcop.com, and eesel.ai, cross-referenced; Jon Dick quote and HubSpot's official framing as reported by SaaStr and resolve247.ai from HubSpot's own announcement. Intercom Fin per-resolution rate (\$0.99) and \$100M+ ARR figure: eesel.ai and SaaStr, respectively. Zendesk's May 18, 2026 resolution-tier restructuring (Assisted Escalation/Contained Resolution free, only Verified Resolution billed): Zendesk's own help documentation. Zendesk overage rates (\$1.20–\$2.00 range) and the January 2026 auto-billing/no-cap change: Not published by Zendesk; third-party teardowns (eesel.ai, Richpanel, getmacha.com, CorePiper, Voiceflow) — presented as a range for exactly this reason. Salesforce Agentforce pricing history (three-model timeline, deal-conversion figures, ARR, adoption rate, Data Cloud cost range): SaaStr ("Salesforce Now Has 3+ Pricing Models") and getmonetizely.com, cross-referenced; Flex Credit action-cost math is Salesforce's own published rate structure as reported by Enterprise Dreamin' and nsiqinfotech.com. The 75%-of-vendors-have-no-systematic-pricing-approach figure: paid.ai research as cited by diginomica.com. Illustrative cost comparisons in this report use stated assumptions and clearly-labeled estimates where vendors do not publish rates; they are not vendor quotes. If you believe a finding is inaccurate, contact [editorial@unvarnishedreviews.com](mailto:editorial@unvarnishedreviews.com) with documentation.