Futurum RESEARCH

The Pricing Split: Why Core Software and AI Can't Share One Commercial Model

Buyers are Rejecting Per-Seat Pricing for Core Software, but are Asking for it with AI

Analyst(s): Keith Kirkpatrick
Publication Date: August 17, 2026
Document #: AIOKK202608

Key Points

- Buyer Data Cuts Against the "Per-Seat Pricing Is Dead" Consensus: Most current pricing commentary treats seat-based billing as collapsing across the board under agentic AI. Futurum's 2H 2026 Enterprise Software Decision-Makers Survey finds buyers are not rejecting per-seat pricing outright. They are splitting it by category, rejecting it for core software while asking for it back specifically for AI.[1]

- Per-Seat Pricing Ranks Last for Core Software and First for AI, Among the Same Buyers in the Same Survey: It finishes at 12.6% for core software, behind consumption-based (28.9%) and outcome-based (22.2%) approaches, and then finishes first at 42.3% for AI functionality billed separately, ahead of consumption (36.6%) and outcomes (21.1%).[1]

- The Split Runs by Company Size as well as by Product Line: Small and very large enterprises are converging on outcome-based AI pricing from opposite directions, while midsized companies are holding onto add-on and consumption models, a segmentation not showing up elsewhere in the current pricing debate.[2]

- Vendors Have Already Split Into Two Camps, and Neither Fully Matches the Buyer Data: Zendesk, Intercom, and Decagon have gone pure outcome-based; Salesforce, Adobe, Pegasystems, and ServiceNow are hedging with hybrid consumption-and-outcome frameworks.[3]

- ETR Buyer-Panel Data Does Not Yet Reward the Pure Outcome-Based Bet With Stronger Spending Momentum: Across four Technology Spending Intentions Surveys (TSISs) (October 2025 through July 2026), Zendesk's and Intercom's Net Score has run flatter or more volatile than most of the hybrid camp. ServiceNow, Salesforce, and

[1] Futurum Research, 2H 2026 Enterprise Software Decision-Makers Survey, 2026.
[2] Futurum Research, Futurum Research 2026: Key Issues and Predictions ("The Value-Based Shortlist Mandate" and "Pricing Optionality to Become Table-Stakes by Q4 2026"), July 2026.
[3] Vendor pricing positioning as reported in Futurum Research, Futurum Research 2026: Key Issues and Predictions, July 2026.

Adobe hold steadier, higher scores, while Pegasystems' Net Score fell sharply in the most recent survey.[4]

- Futurum Expects Pricing Optionality to Become a Shortlist Gate, Not a Negotiating Point: By the fourth quarter of 2026, Futurum expects procurement teams to screen out vendors lacking a consumption or outcome option for core software and a predictable option for AI, the same way they already screen for security certifications.[2]

Recommendations

1. Run Two Pricing Architectures, Not One: Vendors should stop forcing a single commercial model across their portfolio. Core software should follow the market toward consumption or outcome-based structures; AI functionality billed separately should default to per-seat pricing or a capped, budgetable variant, because that is what the same buyers are asking for in the same survey.

2. Instrument Value Metrics Before Renegotiating Structure: Pricing optionality without proof is a discount by another name. Vendors need transparent, auditable value metrics live in the platform before they change commercial terms, regardless of whether the final model is consumption, outcome, or hybrid.

3. Fund Customer Success as a Revenue Function: Outcome and hybrid models hold only if renewal-quarter engagement matches the sales-quarter effort. Vendors that treat customer success as a cost center should expect outcome-based contracts to unwind at renewal.

4. Segment the Commercial Motion by Company Size, Not Just by Product Line: Small and very large enterprises are converging on outcome-based AI pricing from opposite directions, while midsized companies are holding onto add-on and consumption models.[2] A single enterprise price book applied uniformly across segments will misprice the deal in the middle of that range as often as it gets one right.

What You Need to Know

- Seat-Based Licensing the Default Enterprise Software Commercial Model for More Than a Decade: Futurum's 2H 2026 Enterprise Software Decision-Makers Survey shows buyers who name pricing as a top purchase criterion now rank per-seat pricing last among five structures for core software, at 12.6%.[1]

- Same Survey Finds the Opposite Ranking Once Generative AI Is Billed as a Separate Line Item: Buyers want the predictability of per-seat pricing for AI (42.3%) even as they abandon it everywhere else in the stack. That reversal runs against the prevailing view

[4] ETR TSIS company-level metrics for Zendesk (company ID 1241), Intercom (428), Salesforce (133), Adobe (336), Pegasystems (401), and ServiceNow (281), October 2025 through July 2026 TSIS, live ETR MCP queries run August 7, 2026, independent third-party buyer-panel data. Decagon has no TSIS, Edge, APS, or Observatory coverage as of this survey wave.

that agentic AI is finishing off seat-based pricing altogether. The data says buyers want it kept in exactly the one place most commentary assumes it is disappearing the fastest.[1]

- The Split Is Not Just Between Core Software and AI: Futurum finds small and very large enterprises converging on outcome-based AI pricing, while midsized companies hold onto add-on and consumption models, a company-size dimension the core-software-versus-AI split alone does not capture.[2]
- Vendor Responses Are Already Diverging: Zendesk, Intercom, and Decagon have moved to pure outcome-based pricing, while Salesforce, Adobe, Pegasystems, and ServiceNow are building hybrid frameworks that blend consumption and outcome elements.[3]
- ETR TSIS Spending-Intent Data: For six of these seven vendors (Decagon has no ETR coverage), the data shows the pure outcome-based camp running flatter or more volatile Net Score than most of the hybrid camp over the last four surveys, with Pegasystems as a notable exception inside the hybrid camp.[4]
- Futurum's Broader 2H 2026 Key Issues and Predictions: The report finds 70% of enterprise buyers now favor consumption- or outcome-linked models overall, with pure outcome-based preference nearly doubling year over year.[2]
- Futurum's Prediction on Pricing Optionality: This will become table stakes by the end of the fourth quarter of 2026, with vendors unable to offer flexible pricing models losing consideration regardless of product quality.[2]

Analysis

Enterprise buyers have split their preference into two, and that split runs along a line most vendors, and most of the current pricing commentary, have not drawn. The dominant narrative around AI and pricing holds that agentic AI is replacing the seat; that per-seat billing collapses once software does the work a human used to do. Futurum's buyer data says otherwise: the two categories are moving in opposite directions, and AI is the one where buyers are asking for the seat back, not the one where they are abandoning it the fastest.

Buyers Have Split Their Pricing Preference into Two

Futurum's 2H 2026 Enterprise Software Decision-Makers Survey asked buyers who name pricing as a top purchase criterion to rank five commercial models for core software. Per-seat pricing finished last, at 12.6%, trailing consumption-based pricing (28.9%) and outcome-based pricing (22.2%).[1] That result lines up with a decade-long shift away from flat per-user licensing, accelerated now by the arrival of agentic workers who are not tied to a single human seat.

The same survey then asked a narrower group, buyers whose vendors charge separately for generative AI functionality, to rank the same five models for that AI-specific spend. The ranking flips. Per-seat pricing finishes first, at 42.3%, ahead of consumption (36.6%) and outcomes (21.1%).[1] Buyers pursuing consumption and outcome pricing for the rest of their stack still want a predictable, budgetable line item when the spend category is AI.

Table 1: Pricing Model Preference, Core Software vs. Generative AI Add-Ons

Pricing Model	Core Software	GenAI Add-On
Per-seat	12.6%	42.3%
Consumption-based	28.9%	36.6%
Outcome-based	22.2%	21.1%

Source: 2H 2026 Enterprise Software Decision-Makers Survey, Futurum Research, August 2026

Predictability for AI, Flexibility for Everything Else

The split is not buyers being inconsistent. It reflects two different risk calculations running at once. Core software costs are familiar, and buyers have a decade of benchmark data to negotiate consumption or outcome terms with confidence. Generative AI costs are new, volatile, and tied to usage patterns few procurement teams have modeled yet, so the same buyer who will gamble on outcome-based pricing for a CRM platform wants a ceiling on what an unfamiliar AI feature can cost in a bad month.

Vendors selling both categories under one pricing philosophy are answering a question their buyers did not ask the same way twice. A vendor that forces consumption pricing onto its AI add-on to match its core-platform story is offering flexibility where the buyer wants a cap. A vendor that keeps per-seat pricing across its whole portfolio to simplify billing is offering predictability where the buyer wants to pay only for what gets used.

Vendors Are Answering With Two Different Playbooks

The market has already produced two visible responses to this split, though neither fully resolves it yet. Zendesk, Intercom, and Decagon moved early and fully to pure outcome-based pricing, betting that buyers will accept usage risk across the board if the vendor absorbs enough of that risk itself.[3] Salesforce, Adobe, Pegasystems, and ServiceNow took the hybrid route, blending consumption and outcome-linked elements into their commercial frameworks rather than picking one model for the whole portfolio.[3]

Neither camp has fully built the two-tier structure the survey data calls for: consumption or outcome pricing for core software, paired with a separate, capped, or seat-based structure specifically for AI functionality. The vendors closest to that answer are the ones already billing AI as a distinct line item and treating it as a separate pricing conversation from the core platform, rather than folding it into the same negotiated discount structure.

The commercial risk sits on both sides of this choice. A pure outcome-based vendor that has not also solved for AI-specific predictability is asking buyers to accept usage risk twice, once on the core platform and again on the AI layer, at the exact moment those buyers are asking for a cap on one of the two. A hybrid vendor that blends consumption and outcome elements without separating the AI line item is easier to sell today, but it inherits the same mismatch the survey data flags: buyers steering toward flexibility for software and predictability for AI will eventually notice when a single blended structure gives them neither cleanly.

What ETR's Buyer Data Shows About Each Camp

Futurum's own buyer-panel research measures stated preference. ETR's Technology Spending Intentions Survey (TSIS) measures what buyers are actually doing about it: whether the same population of enterprise respondents is increasing, holding, or cutting spending with each of the vendors named above. Six of the seven vendors in this report carry active TSIS coverage; Decagon does not appear in any ETR dataset, so its early move to pure outcome-based pricing has no buyer-panel evidence behind it yet.

ETR's company-level Net Score reflects spending intent across each vendor's full portfolio, not isolated to the specific pricing motion discussed here, so it is a read on overall momentum rather than a direct verdict on any one commercial model. A company's Net Score is the percentage of respondents reporting positive spend trajectory (planning to adopt a vendor's software or increasing spending), minus the share of respondents that are reporting decreasing spending (planning to decrease spending or replace, put in containment, remove or swap-out a vendor). Therefore, higher Net Scores = a positive spend trajectory, whereas lower Net Scores = a flat or negative spend trajectory.

Table 2: ETR TSIS Company-Level Net Score and Pervasion, Oct 2025 to Jul 2026

Vendor	Camp	Oct-25 NS	Jan-26 NS	Apr-26 NS	Jul-26 NS	Jul-26 Pervasion	Jul-26 NS
zendesk	Outcome-based	-0.5%	+4.4%	-1.3%	-1.7%	14.8%	237
INTERCOM	Outcome-based	-29.2%	+18.2%	+21.3%	-11.1%	2.5%	36
salesforce	Hybrid	+19.7%	+25.8%	+21.3%	+19.0%	33.4%	965
Adobe	Hybrid	+16.9%	+18.0%	+15.7%	+15.5%	36.6%	883
PEGA	Hybrid	+3.7%	+17.0%	+12.4%	+0.9%	24.7%	148
servicenow	Hybrid	+42.2%	+44.4%	+41.7%	+42.8%	43.2%	608

Source: ETR TSIS company-level metrics, live ETR MCP queries run August 7, 2026.[4]

The two vendors furthest into pure outcome-based pricing show the least settled Net Score of the group. Zendesk's Net Score has stayed within a few points of zero for four straight surveys, last at -1.7% in the July 2026 TSIS (N=237), with Pervasion drifting down from 17.2% to 14.8% over the same year. Intercom's Net Score has swung from -29.2% to +21.3% and back to -11.1% across the same four surveys on a thin base of 33 to 48 citations per quarter, below the citation count ETR recommends for confident vendor-level reporting. The direction of any single Intercom quarter should be read with caution; the instability itself is still a genuine signal.

Among the hybrid camp, three of four vendors show steadier, positive Net Score momentum. ServiceNow has held in the low-to-mid 40s across all four surveys (+42.8% most recently, N=608), and Salesforce and Adobe have both held stable positive scores in the high teens (+19.0% and +15.5% most recently, N=965 and N=883). Pegasystems is the exception inside its own camp: its Net Score fell from +12.4% in April 2026 to +0.9% in July 2026, the sharpest single-quarter decline among the six vendors tracked here, alongside a Pervasion dip from 27.1% to 24.7% (N=148 to 191, above ETR's minimum reporting threshold but below the level recommended for peer-set ranking).

None of this establishes that a vendor's pricing model is driving its Net Score. Net Score reflects a vendor's full portfolio and moves with product execution, competitive pressure, and macro budget conditions as much as with commercial terms. What the data shows is narrower and still useful: buyer spending intent has not yet rewarded the pure outcome-based bet with clearly stronger momentum than the hybrid bet, and Pegasystems' hybrid position has weakened enough in the most recent survey to warrant its own watch item alongside the outcome-based early movers already flagged below.

The Split Runs by Company Size and Product Line

The core-software-versus-AI split is not the only fault line in the data, and it is not the fault line most pricing commentary is currently drawing, either. Futurum's research finds small and very large enterprises converging on outcome-based AI pricing from opposite directions: small companies because they cannot absorb a fixed AI bill during a slow quarter, large enterprises because they have the usage volume and negotiating power to demand it, and the internal data teams to define and audit the outcome. Companies in the middle of the size range hold onto more traditional add-on or consumption pricing.[2] A vendor selling a single AI price point across all three segments will misprice the deal in the middle of that range about as often as it gets one right.

This matters most for vendors whose customer base spans the full size range, since a sales team pitching one pricing conversation into three different risk postures will lose deals it should win. The commercial motion needs a segment-specific script before it needs a new number on the price sheet.

Figure 1: Why Company Size Bends Pricing Preference

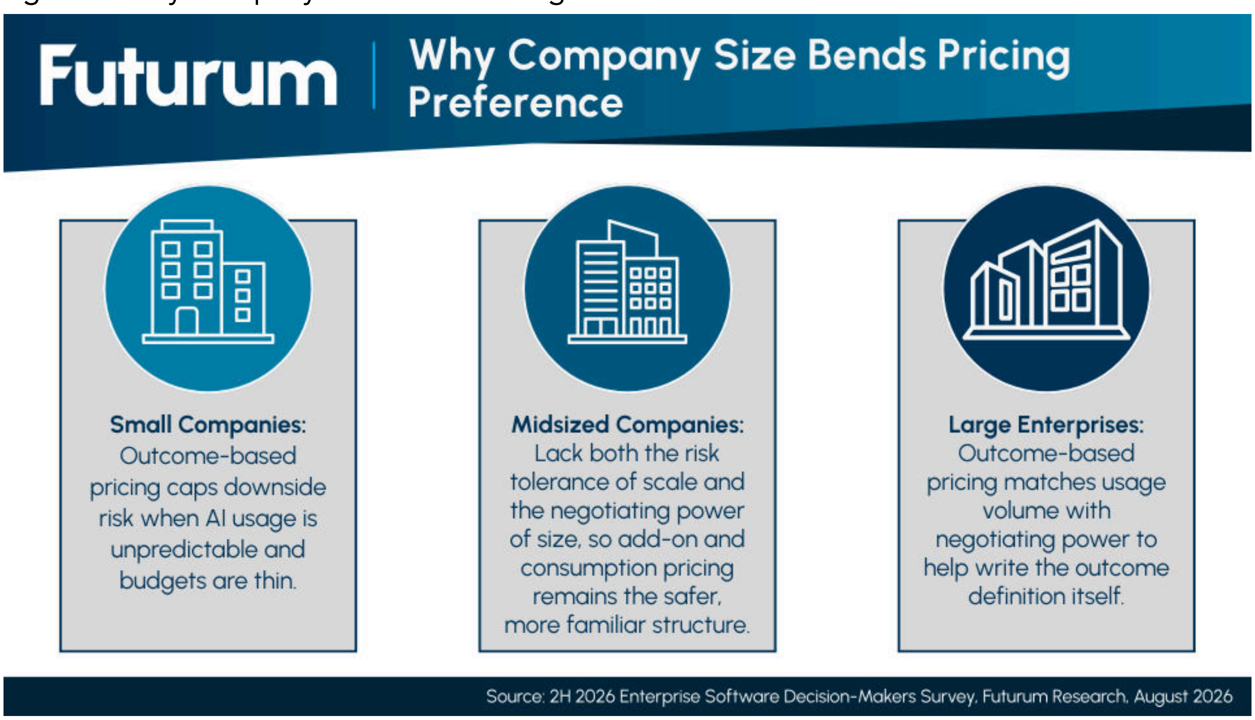

Source: 2H 2026 Enterprise Software Decision-Makers Survey, Futurum Research, August 2026

Future View: A Single Pricing Model Signals Misalignment, Not Simplicity

Futurum's broader research finds 70% of enterprise buyers now favor consumption- or outcome-linked models overall, and pure outcome-based preference has nearly doubled to 27% in a year.[2] Set against a market moving that fast, a vendor holding a single commercial model, whatever that model is, is telling procurement teams something about how well it understands its own buyer, not just about price.

Futurum expects this to matter most at the shortlist stage rather than the final negotiation. Procurement teams filtering vendors down to a shortlist increasingly screen out pricing rigidity before they screen out feature gaps, because rigidity signals a vendor has not adapted its go-to-market to how the category buys today.

By the fourth quarter of 2026, Futurum expects vendors without pricing optionality, meaning at minimum a consumption or outcome option for core software and a predictable option for AI, to be excluded from consideration on commercial architecture alone, the same way vendors already get screened out for missing security certifications.[2]

The fix is not one new pricing model. It is three coordinated moves: instrument the platform to produce transparent value metrics regardless of which structure a given customer chooses, build enough pricing optionality that procurement can select the structure matching its own risk appetite, and fund customer success as an ongoing revenue function rather than a post-sale cost center, since outcome and hybrid deals require continuous engagement to keep the co-owned outcome real.[2]

None of these three moves requires abandoning a vendor's existing pricing philosophy for core software. What they require is treating AI as its own pricing category, with its own predictability requirement, its own value metrics, and its own renewal motion, rather than a feature bolted onto whichever commercial model already covers the rest of the platform. Vendors that make that separation deliberately, before the fourth quarter of 2026 forces the issue, will be negotiating from a position that their competitors are still building.

What to Watch

- Whether the AI Pricing Split Widens or Narrows: Watch subsequent waves of Futurum's Enterprise Software Decision-Makers Survey for whether the gap between core-software and AI-specific pricing preference grows as buyers gain more usage history with AI spend, or narrows as usage becomes as familiar as core-platform costs.
- Renewal Rates and ETR Net Score at the Pure Outcome-Based Early Movers: Watch whether Zendesk's near-zero Net Score and Intercom's volatile Net Score in ETR's TSIS firm up or deteriorate further in subsequent waves, and whether that shift shows up in renewal behavior at either vendor.

- **Whether Pegasystems' Net Score Decline Holds or Reverses:** Watch the October 2026 TSIS for whether Pegasystems' drop to +0.9% Net Score was a single-quarter blip or the start of a trend, since it is the sharpest move among the six vendors with ETR coverage in this brief.
- **Whether Incumbents Formalize a Two-Tier Structure:** Watch whether Salesforce, Adobe, Pegasystems, and ServiceNow move from blended hybrid frameworks toward explicitly separate pricing architectures for core software and AI functionality.
- **Procurement Language in RFPs:** Watch whether enterprise RFPs begin explicitly requiring vendors to offer multiple pricing models as a gating criterion, the same way they already gate on security certifications and service-level agreements.
- **Whether Midsized Companies Hold Their Position:** Watch whether midsized companies continue favoring add-on and consumption pricing as small and very large enterprises converge on outcome-based AI pricing, or whether the middle segment eventually follows one end of the size range rather than the other.

Other Insights from Futurum

Salesforce's Agentic Enterprise Index: A Paradigm Shift in AI Deployment

Is AI Governance Finally Taking Center Stage in Enterprise Strategy?

The Data Foundation Problem: Why "Bring AI to Your Data" Beats the AI Data Lake

About Us

About the Authors

Keith Kirkpatrick is VP & Research Director, Enterprise Software & Digital Workflows for The Futurum Group. Keith has over 25 years of experience in research, marketing, and consulting-based fields. He has authored in-depth reports and market forecast studies covering artificial intelligence, biometrics, data analytics, robotics, high-performance computing, and quantum computing, with a specific focus on the use of these technologies within large enterprise organizations and SMBs. He has also established strong working relationships with the international technology vendor community and is a frequent speaker at industry conferences and events. Keith holds dual Bachelor of Arts degrees in Magazine Journalism and Sociology from Syracuse University.

About The Futurum Group

Every day, The Futurum Group's analysts, researchers, and advisors help business leaders worldwide anticipate tectonic shifts in their industries and leverage disruptive innovation. Unlike traditional analysts, The Futurum Group works not only in analysis and research but also takes that insight and knowledge even further, engaging all the way through the go-to-market process.

Futurum Research provides in-depth research and insights on global technology markets using advisory services, custom research reports, strategic consulting engagements, digital events, go-to-market planning, and message testing. It also creates, distributes, and amplifies rich media content that all stakeholders read, watch, and listen to.

See more details on The Futurum Group at futurumgroup.com.

Copyright & Use License