
W H I T E P A P E R

June 2026

The Five Keys to AI Pricing

Why software companies are getting AI pricing wrong — and leaving money on the table

pricing@blueridgepartners.com
www.blueridgepartners.com

EXECUTIVE SUMMARY

The Five Keys to AI Pricing

AI functionality is now live in most B2B SaaS products, but P&L impact is lagging. In Blue Ridge Partners' recent survey of 128 software CEOs, CFOs, CPOs, and CMOs, **nearly a quarter of companies still bundle AI at no charge**, only **27% have formal, enforced margin guardrails**, and just **8% used willingness-to-pay research as the primary basis for setting AI prices**. The product decision was made with urgency. The pricing decision was made without sufficient rigor.

This white paper lays out a practical operating system for AI monetization, anchored in five interdependent keys. ROI defines the value ceiling. COGS defines the margin risk. Packaging defines what customers can buy. The pricing metric and structure determine how value scales. Willingness-to-pay research calibrates the model to the market. Each depends on the others. None can be executed well in isolation.

Key Findings

- **AI pricing is being repriced under live conditions.** 96% of executives expect to change their AI pricing model and nearly a third want to but have not yet determined how.
- **Margin exposure is largely unmeasured.** Only 23% of companies have linked AI feature costs to customer-level delivery; 90% of CXOs say AI-driven gross margin compression has already surfaced — or will within six months — with their board, investors, or in diligence.
- **Renewals are now where pricing strategy is won or lost.** 60% of vendors have faced meaningful renewal friction tied to AI: 37% have seen pushback on AI-tied price increases, and 24% have had customers argue AI should reduce their bill.
- **Seat-based pricing is structurally exposed.** 57% of vendors know or believe their customers have non-human users already generating value, yet only 22% of those have explicit pricing for it.
- **While variable pricing can protect margins, there are many challenges to making this work for most software companies.** In a recent Blue Ridge Partners buyer survey, 0% of B2B software buyers preferred a token-based pricing metric.

The Five Keys to AI Pricing

- **Quantify credible ROI** — substantiate AI value in customer-specific, dollar-quantified terms.
- **Price to value with margin protection** — hold value-based pricing, with cost structure as the floor.

- **Package AI for adoption** — integrate AI into core Good / Better / Best tiers rather than gating it behind a standalone SKU.
- **Choose the right pricing metric and structure** — pick a metric that correlates with value, is forecastable, and trends upward.
- **Set prices through willingness-to-pay research** — calibrate prices, discounts, and packaging to blinded buyer evidence.

Over half the financial impact from AI repricing comes from existing customers — renewal upsell and cross-sell — not new logos. Capturing it requires ROI evidence, sales enablement, migration paths, and discount guardrails in place before the renewal conversation starts.

The window is closing. Companies that move now can still shape buyer expectations, defend gross margin, and capture AI-driven revenue. Companies that wait will reprice under duress — at renewal when the customer has anchored to the old model, after a competitor has redefined the category, or in diligence when buyers are setting the terms.

Part I. The State of AI Pricing in Software

AI is reshaping software pricing more rapidly than any prior shift in the industry's history, including the move to per-user SaaS pricing in the early 2000s and to consumption pricing in the late 2000s. Buyer expectations of AI ROI have changed. Cost has become variable for the first time in modern software. Packaging conventions are being rewritten in real time. Pricing metrics that worked five years ago no longer correlate with value. Software companies that wait — or that make pricing decisions without fresh market research — are leaving substantial revenue uncaptured.

Blue Ridge Partners' recent survey of 128 SaaS CEOs, CFOs, CPOs, and CMOs found that while most companies have AI live in their product, **25% are still bundling AI into existing subscriptions with no incremental charge**. Among those that do charge, 79% have attached at under 50%, and for 82% of companies with paid AI tiers, AI-specific pricing still represents under 10% of total ARR. The gap between “we have a pricing model” and “AI is showing up in the P&L” is the story most boards are struggling to absorb.

While the product decision was made with urgency, the pricing decision was made with guesswork. AI-native competitors are entering established categories with modern pricing models, no legacy contracts, and no installed base to protect. Incumbents that wait too long will not just miss the revenue — they risk letting others define what AI should cost in their own markets.

The window is closing from several directions at once. Pre-AI contracts have no pricing provisions for AI — every renewal that passes without a strategy locks in bad economics for another two to three years. AI inference costs scale with adoption, meaning margin exposure grows every day usage increases. And non-human identities (NHIs) are growing inside customer environments while pricing models still assume the user is a person in a seat.

Even the largest software companies are still working through it. Salesforce has moved through named-user pricing, per-conversation charges, bundled conversation allotments, and flex pricing — four AI pricing revisions in twenty-four months. A senior Salesforce executive said in May 2025 that “any business that thinks they’ve figured out AI pricing is kidding themselves.” The market has not settled, and most companies do not have Salesforce's room to experiment in public.

Financial Rigor Was Deferred

At most software companies, AI pricing followed a familiar pattern. Product shipped because the market demanded it. Pricing was developed quickly because launch timelines were moving. The harder work — willingness-to-pay research, feature-level cost modeling, margin analysis — was deferred.

The data bears this out. Only 27% of companies have formal, enforced finance guardrails on AI pricing. For the remaining 73%, margin thresholds either do not exist, are not enforced, or are still being

established. 72% priced AI with only informal research or none at all. Only 8% said willingness-to-pay research primarily drove the decision.

That matters because AI changes the economics of software delivery. AI features look like software to the customer, but behave like a variable-cost service to the vendor — every query, workflow, or recommendation creates a cost that scales with usage. Most companies launched AI without knowing what it would cost at the feature level, which customers would use it most, or whether the price would hold up against the margin profile investors expect.

Metric choice compounds the problem. In a recent Blue Ridge buyer survey, 0% of B2B software buyers preferred a token-based pricing metric. Token and usage-based structures remain common among companies that have not revisited their approach since launch. Most vendors chose the metric before asking buyers what they could understand, forecast, and defend to their own CFOs.

Cost Structures Are Largely Unknown

Traditional SaaS economics were built on costs predictable enough to manage. AI inference breaks that. AI cost is variable, grows with usage, and in most companies is not tracked by customer, feature, or workflow.

Only 23% of companies have fully linked AI feature costs to customer-level delivery. The rest are tracking at a product level at best, with no reliable view of which customers are margin-accretive and which are being subsidized. A significant share are renewing contracts without knowing the answer.

Executives need to be prepared to prove at an account level which features are profitable and which customers are being subsidized. 90% say gross margin compression from AI costs has come up — or will — with their board, investors, or in a fundraising or M&A process.

Renewal Is Where Pricing Strategy Becomes Financial Reality

The renewal cycle is where deferred pricing decisions stop being theoretical. Companies with pre-AI contracts are delivering AI functionality and absorbing inference costs with nothing in the contract to recover them. The problem is compounding with every cycle.

37% of vendors have faced customers pushing back on AI-tied price increases, and 24% have encountered customers who argued that AI should reduce their bill, not raise it. Combined, **60% have faced meaningful renewal friction**. Without a clear strategy, sales teams will discount to protect the relationship, migration plans will stall, and the company will defer the same problem for another term.

In our experience, over half of the financial impact from AI pricing changes comes from renewal upsell and cross-sell programs rather than new logos. Capturing it requires ROI evidence, sales enablement, migration paths, and discount guardrails in place *before* the renewal conversation starts, not during it.

The Non-Human Identity Problem Is Structural

Non-human identities (NHIs), such as AI agents, automated workflows, service accounts and bots, already outnumber human users 50 to 1 in many enterprise environments. Seat-based pricing was built for a world where value tracked reasonably with human users. That world is changing faster than most pricing models.

57% of vendors know or believe their customers already have non-human users generating value from their software. Of those, only 22% have explicit pricing for it — and half have no pricing at all, either absorbing NHI usage under seat-based contracts or still evaluating what to do. 74% say monetizing that usage is highly important to revenue and valuation. Procurement teams are already arguing that automation should reduce seat counts and bills. Most vendors are not ready to respond.

Companies that define a credible approach now will capture this usage as it scales. The ones that wait will find procurement has written the answer first.

Part II. The Five Keys to AI Pricing

Most CEOs and CFOs recognize that today's AI pricing models are temporary. In our survey, **96% expect to change their AI pricing model and nearly a third want to but have not yet determined how.**

Companies need to make these changes before the market, customers, and investors force the issue.

Based on our work with hundreds of software companies, we see five requirements that determine whether AI pricing shows up in the P&L or gets left on the table. They are interdependent. Pricing without ROI evidence weakens the value case. Pricing without COGS visibility creates margin risk. Packaging without a clear migration path slows adoption. Choosing the wrong metric or structure leaves future usage unmonetized.

Done well, the five keys work as a system

1. **Quantify credible ROI.** Buyers are no longer accepting generic AI claims. They want evidence that AI improves revenue, productivity, risk reduction, or cost efficiency. Companies that can quantify AI impact in customer-specific terms will justify higher prices, reduce renewal friction, and win deals faster.
2. **Price to value with margin protection.** AI has introduced real variable costs into software economics. Companies need to understand their costs by customer, segment, and feature, and use that insight to protect gross margin. The goal is value-based pricing with cost structure as the floor — not cost-plus pricing that caps upside or flat pricing that subsidizes the heaviest users.
3. **Package AI for adoption.** AI packaging should make it easy for customers to buy, expand, and migrate. The highest-ROI, most differentiated AI capabilities should support premium-tier placement or expansion motions. More common AI functionality should raise the value of the core product and support entry-price increases.
4. **Choose the right pricing metric and structure.** The right metric must correlate with value, be easy for customers to understand and forecast, and continue to grow as AI usage scales. Seat-based models must evolve to monetize non-human identities. For many B2B SaaS companies, the best answer will be a hybrid structure — a platform fee plus a value-correlated variable that captures upside and protects predictability.
5. **Set prices through willingness-to-pay research.** AI pricing cannot be set confidently through internal debate, stale benchmarks, or sales feedback. Blinded buyer research — in which respondents do not know who commissioned the study — determines optimal list prices by segment, establishes discount guardrails, validates metric and packaging choices, and surfaces where revenue is being left on the table.

Together, the five keys are the operating system for AI monetization. ROI defines the value ceiling. COGS defines the margin risk. Packaging defines what customers can

buy. The metric and structure determine how value scales. Willingness-to-pay research calibrates the model to the market.

Key 1. Quantify Credible ROI

Buyers' expectations for AI ROI have changed faster than most software companies have adapted their pricing. Two forces are driving the shift. First, headline-grabbing enterprise cost-savings stories have made buyers expect concrete, dollar-quantified outcomes. IBM, for example, has publicly attributed roughly \$3.5 billion in productivity gains to AI agents and projected that nearly 8,000 back-office roles will be replaceable by AI over five years¹.

Second, the wave of inflated and unsubstantiated ROI claims that accompanied last year's rush of AI product launches has made those same buyers significantly more skeptical of vendor claims. A credible ROI claim — one the buyer can stress-test — justifies higher prices, accelerates sales cycles, and lifts win rates. An overstated claim risks disqualifying vendors from serious consideration.

The ROI from agentic or other AI functionality can be powerful and deliver game-changing impact. One SaaS CEO we worked with confirmed our analysis: “Our new agentic AI functionality fully doubles the ROI impact of our software, from X% EBITDA uplift to 2X%.” Substantiating the additional financial impact of AI functionality can critically enhance monetization and acceleration.

The strongest path to credibility is partnering with early customers to validate realized ROI and securing their willingness to act as references. Where that data is not yet available, commissioning an independent third-party ROI study, and then confirming through buyer research that the described functionality would actually deliver it, can be equally effective. Both approaches can be packaged as sales and marketing assets and materially shift the renewal negotiation.

ROI research also unlocks value well beyond justifying price. Not all AI functionality or use cases deliver the same return. Understanding which features and use cases generate the largest customer outcomes should inform product strategy, sharpen sales targeting, and could become the foundation for an outcome-based pricing structure. Software companies should ensure they capture usage telemetry to substantiate those ROI claims over time. The same telemetry that substantiates ROI claims to a prospect can also support customer success management, stabilize retention, and enable expansion at renewal.

Best-practice recommendations

- Perform thorough ROI research on your product's AI functionality and incorporate credible ROI evidence into sales and marketing
- Partner with early customers to validate realized ROI and secure references
- Capture the product telemetry needed to substantiate ROI claims over time

¹ IBM Think Insights, “Embracing the future of HR by becoming an AI-first enterprise”; Arvind Krishna remarks, Bloomberg News, May 2023.

Key 2. Price to Value with Margin Protection

The new wrinkle with AI is that cost now constrains pricing for the first time in modern SaaS. Traditional software economics worked because COGS — hosting, support, and infrastructure — was largely fixed and amortizable, which is why per-seat pricing scaled cleanly. AI inference (the per-use cost of running a prompt through an AI model) flips that. Inference cost is variable, marginal, and directly correlated with usage. The result is the first time most software companies have had to factor real product costs into how they set prices.

The reflex of many CFOs is to switch to usage-based pricing, so that revenue moves with cost and gross margin is protected. The instinct is rational, but two problems undermine it in B2B SaaS. First, customers cannot reliably forecast their own usage, which kills their ability to budget and suppresses willingness-to-pay. Second, usage and value frequently diverge. A customer using 2x the tokens does not necessarily realize 2x the value. Pricing designed around cost drivers rather than value drivers leaves money on the table.

Why, then, does usage-based pricing work for the four frontier model providers — Anthropic, OpenAI, Google, and Microsoft (via Azure OpenAI)? These companies serve as industry utilities, much like AWS and Azure before them. At the API layer, usage *is* the value, buyers are technical, and procurement organizations have built muscle around hyperscaler-style consumption pricing. None of those conditions reliably hold in B2B SaaS, where buyers are business users, value is delivered through workflows rather than tokens, and the procurement organization is built around predictable line items.

The biggest practical challenge in factoring COGS into software pricing is understanding what inference costs will actually be across different customer sizes, segments, and use cases. The best approach is to run instrumented pilots with early customers and use the resulting data to build an empirical view of cost drivers. Without this baseline, pricing decisions are guesses.

Critically, pricing should not be “cost-plus,” which structurally caps the price at the level of a competent reseller and leaves the entire value premium uncaptured. COGS belongs in the analysis as a floor and as a risk-management input to size bundles and overage thresholds — not as the basis from which price is built up.

One effective construct we see today is a value-based price that includes a meaningful baseline of usage in each tier (e.g., “Tier 5 includes up to X”) with overage charges beyond it, structured like a mobile phone data plan. This protects margin against the long-tail power user without forcing every customer to forecast their own consumption.

One final, often-missed point: how usage is defined matters as much as how it is priced. B2B buyers have strong, consistent resistance to token-based metrics they do not understand. In one Blue Ridge buyer survey, 0% of respondents preferred a token-based usage metric for purchasing B2B software. Usage should be defined in terms tied to value — actions that produce a customer outcome (e.g., tickets

resolved, documents generated, qualified leads identified, contracts reviewed). The metric should be intuitive, auditable, and clearly connected to what the buyer is paying for.

Best-practice recommendations

- Analyze LLM inference costs and correlate them with usage patterns by customer type, segment, and use case
- Define usage in customer-meaningful terms (e.g., tickets resolved, documents generated) — not in tokens
- Protect margin with bundled-usage tiers plus defined overage charges
- Monitor inference cost monthly, in aggregate and by segment, to tune bundle sizes and overage thresholds as usage patterns evolve

Key 3. Package AI for Adoption

Packaging decisions translate the value framework from Key 1 into a structure customers can buy. The central question for AI features is whether they should be sold as a standalone bundle or integrated across existing Good / Better / Best tiers. Many software companies have answered this question one way in the short term and will need to answer it differently in the long term.

Today, a number of leading vendors, including Microsoft (M365 Copilot), Salesforce (Einstein add-ons), HubSpot (Breeze AI), and Atlassian (Atlassian Intelligence), package AI as a separate, premium SKU sold alongside the core platform. The logic is sound for now. AI is still differentiated, buyers perceive it as new, and a discrete SKU captures the willingness-to-pay premium that novelty creates.

That window is closing quickly. As AI becomes embedded across product categories, a non-AI bundle becomes a structurally inferior product. At that point, a separate AI SKU shifts from differentiator to forced and awkward upsell — a drag on win rates and renewal motions.

The better long-term answer is integrating AI directly into core Good / Better / Best bundles, with tiering driven by relative value and competitive differentiation (leveraging the ROI hierarchy from Key 1). Not all AI functionality delivers the same return. Mapping features to their realized customer outcomes — from commodity-level productivity gains to high-value workflow automation and revenue impact — creates a hierarchy that informs both product strategy and packaging. The highest-ROI features justify premium-tier placement; the lowest become table-stakes inclusions in the entry tier.

There are two things to watch for as you design bundles. First, packaging cascades directly into unit economics: putting inference-heavy features into the entry tier without usage caps is a fast path to compressed gross margin. Packaging design and COGS design must be done together. Second, visible AI (e.g., copilots, generative tools) belongs as a differentiator between tiers; embedded AI, such as smart routing, classification, and anomaly detection, is best treated as a product-wide quality uplift rather than a tiering lever.

Best-practice recommendations


- Plan for AI to be integrated across core product Good / Better / Best bundles — not as a separate AI add-on bundle, even if a separate SKU is the right transitional play today
- Use value, differentiation, and COGS to decide which AI capabilities go in which tier: commoditized features anchor the entry tier; differentiated, high-ROI features earn premium-tier placement
- Coordinate packaging design with the COGS analysis from Key 2 — bundle inference-heavy features only with usage caps sufficient to defend gross margin
- Distinguish visible AI from embedded AI: treat the first as a tiering lever and the second as a product-wide quality uplift

Key 4. Choose the Right Pricing Metric and Structure

The shift in pricing metrics underway today is happening faster than any prior pricing transformation in the industry's history. Previous comparable transitions — the shift to user-based pricing with the rise of SaaS in the early 2000s and to consumption-based pricing in the late 2000s as AWS legitimized the utility model — each unfolded over a decade. The current AI-driven wave is forcing more software companies to revisit their pricing metric in 12–24 months than either of those prior migrations did over their full course.

Three forces are driving the shift. First, user-based pricing increasingly fails the value test, because AI agents and automated workflows now perform work previously done by users. A customer who replaces five seats with one agent should not see their bill drop by 80% if value is unchanged or higher. Second, customers are more skeptical of pricing metrics that do not track to value or that they cannot forecast and budget. Third, the visibility of AI ROI (Key 1) is creating real interest in outcome-based pricing, where price moves directly with realized customer value.

Figure 1. *The six pricing metrics most commonly used by software companies — relative usage and trend direction*



Outcome-based pricing, such as charging per resolved support ticket, has earned attention because it correlates closely with value in concept. In practice, however, our experience shows it underperforms for the majority of software companies. Buyers and vendors often cannot agree on the definition of the outcome, buyers do not give the software full attribution for the outcome, and most problematically, the outcome itself does not always correlate with willingness-to-pay. Companies considering outcome-based models should pressure-test all three before committing.

Correlation with willingness-to-pay is the single most important test of any pricing metric. When a metric correlates well, buyers are willing to pay materially more as their use of the metric grows: they accept that 10x the users, 10x the transactions, or 10x the resolved tickets is worth meaningfully more. When correlation is weak, the same metric will produce prices that are dramatically too high for some customers and too low for others, which undermines pricing credibility, depresses win rates, and leaves revenue on the table. Validating correlation requires willingness-to-pay survey data, covered in Key 5.

The best-fit pricing metric satisfies five criteria

- Correlates with value, both intuitively and based on quantitative willingness-to-pay analysis
- Future-proof: the metric trends upward over time, not downward
- Easy for the customer to forecast and budget against
- Easy to track and audit by both the vendor and the customer
- Simple to understand and communicate to the buyer

Pricing Metric	Examples	Trend & Challenges	Today's AI Environment
User-Based	Price per named user; price per concurrent user (Salesforce, HubSpot, Microsoft 365)	Declining over the past 10 years because user count does not always correlate with full value, and customers can reduce users to reduce cost	Less future-proof as AI agents replace human seats; difficult to monetize agentic users on a per-seat basis
Consumption-Based	Price per email sent; per API call; per token or inference (Twilio, Snowflake, AWS, Anthropic, OpenAI)	Growing for 10+ years on the AWS/Azure precedent. Often hard for customers to forecast; does not always correlate with value	Strong buyer resistance to token-based pricing in B2B SaaS. Works at the frontier-model layer; rarely at the application layer
Fixed Fee	Annual flat fee for unlimited consumption ("all-you-can-eat")	Hard to communicate solution value and tie value to price. Margin risk under heavy users	Difficult to absorb variable LLM inference costs; viable primarily as the platform-fee component of a hybrid structure
Firmographic-Based	Price per employee; per revenue or AUM tier (HR tech, fintech, vertical SaaS)	Declining where the firmographic does not correlate tightly with realized value	Workable for AI-enabled products where the firmographic correlates with value; usually pairs with a hybrid structure

Pricing Metric	Examples	Trend & Challenges	Today's AI Environment
Outcome-Based	Per resolved ticket (Intercom Fin, Salesforce Agentforce); per qualified meeting; per cyber threat remediated; % of recovered receivables	Emerging rapidly. Requires buyer-vendor agreement on outcome definition and attribution. Real risk of willingness-to-pay misalignment	Widely being evaluated for AI features. May not satisfy all five metric criteria, particularly forecastability and willingness-to-pay correlation
Hybrid	Platform fee + variable metric (Snowflake compute + storage; Salesforce platform + Agentforce per-conversation; Twilio platform + per-message)	Rising in adoption, particularly among mid-to-large vendors. Captures usage/outcome upside while preserving budget predictability	Strongest fit for AI-enabled software. Mitigates the weaknesses of any single metric, but requires careful research to size each component

A note on structure

Selecting the right metric is half the work; the structure built around it is the other half. Typical structures combine a fixed platform component — which provides budget predictability, anchors a relationship-level price, and supports cross-sell — with a variable component tied to a value-correlated metric, which captures upside as customer use grows and protects margin against rising inference costs. Tiered bundled-usage with overage protection, the model discussed in Key 2, is the most common form. Annual commitments with usage true-ups, ramp deals for net-new accounts, and outcome-tied accelerators are useful structural levers depending on the buyer profile.

Best-practice recommendations

- Evaluate multiple metrics across the five criteria above
- Validate willingness-to-pay correlation quantitatively before committing to any metric
- Confirm the metric trends upward over time — avoid metrics declining structurally (e.g., human seat counts, where AI is reducing demand)
- Stress-test outcome-based pricing against all three failure modes (definition, attribution, and willingness-to-pay correlation) before deploying it

Key 5. Set Prices Through Willingness-to-Pay Research


The previous four keys describe how to determine the value of AI functionality and design the packaging and pricing structure around it. This key focuses on setting the right prices to capture that value.

Ask most software CEOs and CFOs whether their sales teams are capturing the full revenue value of their solutions, and the answer is invariably no. Updating pricing using the process described in this paper can produce tremendous improvement in capturing value to the P&L.

Software pricing has always been more dependent on fresh market research than most industries because of the speed of innovation and the intensity of competition. That dependence is even sharper today. Buyer expectations of AI ROI are moving in real time, competitive offerings shift every quarter, and the willingness-to-pay correlations that held a year ago are no longer reliable. Pricing decisions made on stale internal data, win-loss intuition, or internal discussions are dangerously imprecise in this environment.

Setting list prices and discount levels with confidence requires both a blinded buyer survey across target segments — in which respondents do not know which vendor sponsored the research, materially reducing social-desirability bias and inflated willingness-to-pay responses — and a quantitative willingness-to-pay analysis on the survey data, using multiple analytical methods (e.g., conjoint analysis for feature trade-off and packaging, Gabor-Granger for specific price-point validation).

Figure 2. **Example of optimal pricing and discounting ranges determined for a single bundle in a single market segment**



A well-designed willingness-to-pay study answers a specific set of questions that determine pricing outcomes:

- The optimal **list price** by segment and bundle
- The **discount guardrails** below which deal economics break
- The **right pricing metric** — validating or invalidating the Key 4 choice with quantitative evidence on willingness-to-pay correlation
- **Packaging boundaries** — which AI features customers will pay separately for, and which they expect bundled
- **Segment-level differentiation** — buyer cohorts with materially different willingness-to-pay that should see different prices, bundles, or sales motions

Cadence is increasingly important. In a stable market, a foundational pricing study every two to three years was sufficient. In the current AI environment, **foundational research annually, paired with continuous monitoring of win-loss data, discount patterns and cost analysis, is the right baseline.** Companies waiting longer between studies are essentially repricing on stale assumptions.

Best-practice recommendations

- Conduct a blinded buyer survey across target segments and apply willingness-to-pay analysis before setting or refreshing list prices, discount guardrails, or pricing metrics
- Use the willingness-to-pay findings to validate the four prior decisions: ROI claims (Key 1), COGS-aware bundling (Key 2), packaging boundaries (Key 3), and the chosen pricing metric (Key 4)
- Share subsets of the market research with the sales team to strengthen their ability to sell and capture full value
- Refresh foundational research annually in the current AI environment, and treat win-loss data as continuous monitoring between full studies

These elements are interdependent. Credible ROI evidence sets the value ceiling; cost sets the floor; packaging structures the customer's choice; the pricing metric determines how value is captured as the customer grows; and willingness-to-pay analysis is what calibrates each of those decisions to the actual market. None of the first four can be executed well without the fifth.

Part III. Capturing P&L and Valuation Impact

Well-executed AI repricing delivers a material ARR uplift and a meaningfully improved gross-margin story — the one investors and acquirers are now asking about in diligence, board meetings, and fundraising conversations. The companies that capture the prize share three things: (i) they know what their AI is costing at the per-customer and per-feature level; (ii) their sales teams can defend and sell the full value of AI; and (iii) their board can connect AI pricing and monetization to the valuation multiple before anyone asks.

Most companies have none of the three. A significant portion of the planned financial impact from AI pricing programs is not going to be captured — not because the strategy was wrong, but because the execution infrastructure was never built.

This section addresses the gap between price-setting and price-getting — the commercial enablement work that determines whether the decisions made in the previous sections show up in invoiced revenue and improved valuation.

Across our clients, only 37% of pricing-program impact comes from new-logo price increases. Over 50% comes from renewal upsell and cross-sell programs that carry the pricing strategy into existing-customer motions. Companies focused exclusively on new-logo list prices will leave the majority of available impact on the table before they start.

Figure 3. % of Total Pricing Impact by Pricing Enablement Driver for Software/SaaS Companies (source: Blue Ridge Partners)

% of Total Pricing Impact by Pricing Enablement Driver for Software/SaaS Companies*
(% of total pricing-driven revenue growth)



Product Telemetry Is the Foundation

Without knowing what customers are consuming at the feature level, a company cannot reprice rationally, invoice accurately, defend margin in a board conversation, or answer a diligence question about AI gross margin.

Product telemetry is the LLM inference costs tracked by customer, by feature, and by workflow, linked to revenue so that margin can be calculated at the account level. It addresses three problems simultaneously:

- It identifies the least profitable customers — the ones with heavy AI usage against prices set before anyone understood the cost structure, quietly compressing gross margin while NRR looks fine
- It provides the usage distribution needed to design consumption tiers and overage thresholds that protect margin without overcharging the median customer
- It builds the invoicing infrastructure to charge for what is actually delivered as pricing structures evolve toward consumption and outcome models

For companies that launched AI without this instrumentation, usage logs and LLM cost rates can be used to model it retrospectively — not a clean dataset, but sufficient to reprice. The priority is to have telemetry live before the next renewal cycle.

The Sales Motion Has to Change — and That's Where the Money Is

Most sales teams were built to facilitate feature-benefit conversations. The AI repricing conversation is different. Sales teams need to connect functionality to customer outcomes in financial terms, defend a price increase with data, and manage renewal without discounting away the margin the pricing strategy was designed to protect. The teams that can do this capture the uplift; the ones that cannot give it back at the negotiation table.

In the first half of 2026, securing renewal price increases has become measurably harder. Procurement teams have reset expectations, opening with demands for flat or reduced pricing, shorter terms, and more aggressive SLAs. This is partly negotiating posture; our most recent buyer surveys show customers still expect to pay price increases but are demanding decreases as an opening move. It nonetheless means renewal discussions now require commercial negotiation skills that most Customer Success teams are not built for.

Best-practice renewal programs have four elements

1. **Value-based positioning** that explains the price increase anchored on ROI evidence. Without this, sellers default to apologizing for the increase — the most common cause of leakage
2. **Market-research-derived renewal pricing guardrails** that set minimum, expected, and maximum renewal price-increase levels by segment, validated by current willingness-to-pay analysis — the same blinded buyer surveys described in Key 5, expanded to include current customers
3. **Cohort-tailored upsell and cross-sell recommendations for each customer**, pre-mapped before the negotiation starts, so the team is presenting expansion options rather than just defending the renewal price

4. **Sequenced sales plays** — a pre-defined Play 1, Play 2, and Play 3 for each cohort, with clear escalation paths so that the outcome in aggregate is determined by the program, not by individual judgment under pressure

Cross-sell programs operate on the same logic but are not timed to a renewal event. Nearly all of the inputs required already exist if the pricing work in Part II has been completed — the cohort segmentation, willingness-to-pay curves, bundle definitions, and value-selling positioning all transfer directly. The remaining work is operational: building account-level plans that pre-map two or three tailored cross-sell offers per customer at prices already validated against the willingness-to-pay analysis.

Build the Board Narrative Before You Need It

The monetization and valuation conversation around AI pricing is arriving for every software company — in the next board meeting, the next round of fundraising, or the next M&A process. CEOs and CFOs who have built the narrative before they need it will be valued differently from those who get asked the question cold.

An effective board narrative has three components: (i) gross margin visibility; (ii) a credible repricing roadmap; and (iii) a structural repricing strategy that monetizes both human and non-human usage.

Gross margin visibility. Investors need to be satisfied with gross margin visibility at the customer level and, ideally, the feature level. 91% of CEOs and CFOs say that gross margin compression has come up already or is expected to come up. Executives need to show that AI functionality is margin-accretive, that the cost structure is understood and managed, and that NRR alone is not the whole story. NRR can look healthy while AI inference costs quietly compress gross margin inside specific accounts.

Credible repricing roadmap. Investors will ask when and how AI pricing will be updated, what the ARR uplift assumption is, what the migration risk looks like, and what the commercial enablement program will be to ensure good execution. The uplift from well-executed repricing needs to be in the model with data to support it and meaningful milestones — not solely as an aspiration.

Monetizing human and non-human usage. As non-human identities replace human users, investors will ask whether seat-based pricing models have a structural revenue ceiling. 74% of CEOs and CFOs say that monetizing non-human users is critically important — and the challenge is compounding.

The Window Is 6–12 Months

The repricing window exists because the market is still forming. Buyers do not yet have fixed expectations of what AI should cost. Competitors have not established a pricing standard that locks everyone else in. Contracts are coming up for renewal whether companies are ready or not.

The companies that move first will set reference prices and models. Contracts renewed without AI provisions will reset for another cycle. Investors who are currently learning what questions to ask about AI margin will have answers — and will price the gap accordingly.

Act now with leverage, or reprice under duress: at renewal when the customer has anchored to the old model, after a competitor has made the case that your pricing model belongs to a previous era, or in diligence when buyers are setting the terms.

Best-practice recommendations

- Complete the five keys of AI pricing to reset your pricing program and price AI features to value
- Implement product telemetry before the next renewal cycle and use it to continuously improve your pricing program
- Staff the renewal function for commercial negotiation, not just relationship management – CS teams are not equipped for organized procurement
- Build a structured renewal program with all four elements: value positioning, willingness-to-pay-derived guardrails, cohort-tailored upsell plays, and sequenced fallback offers
- Build a parallel cross-sell program owned by core sales, with account-level plans pre-mapped at willingness-to-pay-validated prices
- Construct the board narrative now: feature-level margin visibility, a repricing roadmap with ARR uplift assumptions, and repricing strategy that monetizes both human and non-human usage.

Conclusion

This white paper has outlined an integrated framework for pricing AI-enabled software: substantiate ROI, value-price with margin protection, design packaging for adoption, choose the right metric, validate everything through willingness-to-pay research, and build the commercial enablement programs to capture the value in the field. The decisions made in the next twelve months will shape unit economics, valuation multiples, and exit outcomes for software companies for years beyond that.

About Blue Ridge Partners

At Blue Ridge Partners, we focus exclusively on helping companies accelerate profitable revenue growth — the #1 driver of value creation. We've earned high marks from CEOs, CFOs and private equity firms by delivering material, rapid, and cost-effective results for their portfolio companies.

Our work spans value creation planning, commercial model transitions, strategic pricing, sales effectiveness, and commercial due diligence. Our Managing Directors combine top-tier strategy consulting and operating experience, and we're known for rolling up our sleeves, taking a pragmatic, hands-on approach, and focusing on the “how” of execution to drive measurable P&L impact.

Increasingly, we help clients leverage Commercial AI to accelerate growth. Through our Commercial AI Center of Excellence (CoE), we guide AI investments toward real business problems and commercial metrics — not just tools — to ensure tangible results.

Since 2002, we have partnered with over 130 top-tier PE firms and 1,500 companies worldwide to create lasting value across the investment lifecycle.

To learn more

If AI pricing is becoming a board-level, investor, or renewal priority in your organization, Blue Ridge Partners has helped more than 100 software companies navigate these challenges. Please contact us at Pricing@blueridgepartners.com or visit www.blueridgepartners.com.