

TECHNOLOGY, MEDIA, AND TELECOMMUNICATIONS

Rethinking B2B Software Pricing in the Agentic AI Era

By Naila Dharani, [John Pineda](#), Angel Andeyro, Nate Hamming, Yerim Kim, [Nipun Misra](#), and [Akash Bhatia](#)

ARTICLE AUGUST 13, 2025

The rapid rise of [AI agents](#) is redefining how B2B software vendors capture value. Over the past decade, [pricing strategies](#) have evolved significantly—from perpetual licenses to subscription models, and increasingly toward consumption-based pricing. The emergence of AI agents has accelerated this transformation, driving customers to demand an even tighter connection between price and the actual value delivered. Simply put, more customers want to pay solely for the measurable outcomes their software delivers.

Although outcome-based pricing is a simple concept, it presents a host of challenges. A recent Andreessen Horowitz survey highlighted significant barriers: 47% of buyers struggle to define clear, measurable outcomes; 36% worry about cost predictability; 25% face difficulty aligning on value attribution with vendors; and 24% acknowledge that outcomes often depend on factors outside of vendors' control.

Despite the obstacles, outcome-based pricing is disrupting conventional thinking about seat-based subscription licenses. According to BCG's IT buyers survey, 40% of buyers cite seat reduction as their primary lever to decrease software spending—a shift that the adoption of agentic AI facilitates.

In this new era, software vendors must not only determine which types of agentic AI capabilities to offer but also make two critical pricing-related decisions. First, they need to decide how to package AI agents and broader AI-powered innovations: incrementally, as additional features on top of existing solutions, or included in current offerings. Second, they must figure out how to transition from user-based subscriptions to pricing models better aligned with value delivered, while simultaneously mitigating the risk of revenue loss.

Packaging Agentic Offers: Consider Value and Willingness to Pay

Vendors make their packaging decisions against the backdrop of robust and growing AI investments across industries. In BCG's IT buyers survey, 48% of respondents said they planned to increase their spending on AI and GenAI during the next 12 months. Moreover, customers have demonstrated a willingness to pay a premium for agentic AI features: a 2025 ICONIQ Capital survey highlights that 68% of vendors charge separately for AI enhancements or include these features exclusively within premium tiers.

However, customers are only willing to pay for specific kinds of value. We find that customers perceive greater value in expertise (such as domain-specific insights or proprietary data) and in the complete execution of tasks (such as customer support resolution or end-to-end job booking). Conversely, they show less enthusiasm for commoditized agentic functionalities that handle isolated workflow steps, such as issue identification or automated email quote requests.


For packaging agentic AI solutions, the message is clear: vendors providing AI capabilities that neither complete a job nor add incremental differentiated value should incorporate these capabilities into core offerings rather than selling them separately. Even for more differentiated

capabilities, we generally recommend “seeding” limited access to new agentic capabilities in the base product to encourage customer exploration and demonstrate value before monetization.

When integrating agentic functions into primary products, vendors must carefully analyze the margin implications. Unlike traditional software, the cost structure for AI-driven solutions can be highly variable and expensive. For example, one vendor of customer engagement software has experienced margin variance exceeding 70 percentage points across different customer accounts.

Although cost-plus pricing may seem like a logical response to this margin variability, it introduces two major challenges. First, it undermines predictability for both vendors and customers. Second, as model costs decline over time—a common trend in AI—revenues may erode in parallel unless usage grows exponentially. (See Exhibit 1.)

EXHIBIT 1
Variable AI Costs Turn Cost-Plus Pricing into a Revenue Rollercoaster



Sources: Publicly available model pricing; Internet Archive; Stanford University Center for Research on Foundation Models; BCG analysis.
Note: Leading-edge models are classified based on rankings from the Stanford Center for Research on Foundation Models as of September 2024. The timeline includes top models from each generation—such as GPT-3, GPT-4, and Claude 3 Opus—while excluding midcycle optimizations like GPT-3.5 Turbo or Claude 3.5 Sonnet. The same-generation model depicted is GPT-3.
¹Revenue derived from cost plus a margin.

We expect some vendors to address margin pressure by raising prices for core products—even within existing pricing structures such as per-seat licenses. To mitigate churn risk, they may restrict agentic capabilities to premium tiers, implement usage caps, or strategically position AI enhancements as differentiated and value creating so that they warrant a higher price.

Agentic Pricing Models: Emerging Models and Tradeoffs

As agentic AI offerings become increasingly sophisticated, we expect more vendors to shift from replacing existing software capabilities to providing incremental and value-adding functionalities. These functionalities might include entirely new capabilities previously unattainable, such as accounting software that helps close the books daily or weekly instead of monthly. Alternatively, they may involve tasks that employees previously performed using software, such as recruitment agents that can autonomously review resumes, conduct outreach, and manage candidate screenings.

We have identified five emerging models for pricing these agentic solutions. (See Exhibit 2.) Each one carries specific tradeoffs for both software vendors and customers.

- **Usage-Based: Resources.** Customers pay based purely on their actual use of resources, akin to current cloud computing models. Although this approach provides maximum flexibility for customers, it is also the least predictable. Vendors currently favor it because it mitigates margin pressure stemming from the fluctuating costs of large language models. This method also sidesteps the complexity of defining and measuring the outcomes or interactions that drive value. However, as customers move from pilots to scaled deployments, we expect this model to lose favor. At scale, buyers are more likely to demand cost predictability and clearer alignment between price and delivered value.
- **Agent-Based.** In this model, customers purchase individual AI agents, through either a one-time fee or a subscription. For instance, OpenAI is rumored to be pricing its new PhD-level research agent at \$20,000 per month, mirroring a salary-based cost structure. This model provides customers with clarity and predictability, making it straightforward for companies to reallocate budgets traditionally earmarked for human labor to digital labor.











For vendors, agent-based pricing is most effective when selling multiple agents (by either type or volume) or when agents fully replace human roles. However, complete substitution remains rare. In our survey of IT buyers, 91% expressed interest only in partially autonomous agents, suggesting that customers do not yet view agents as complete replacements for employees. Further, since many use cases and workflows today involve human-agent collaboration, vendors may encounter a ceiling on customers' willingness to pay with this model.

- **Usage-Based: Interactions.** This approach charges customers per defined interaction or activity, from API calls and output generation to human-agent or even agent-agent interactions. It addresses customer concerns identified in the Andreessen Horowitz research by defining measurable outcomes, ensuring clarity in value attribution, and avoiding situations where vendors might be responsible for outcomes beyond their direct control. However, this

model is effective only if vendors can define specific value-adding activities that their agents will perform.

- Outcome-Based: Jobs Completed.** In this model, payment occurs only after AI agents successfully execute specific, predefined jobs, usually combining multiple tasks. It is optimal for vendors with established use cases for which they can guarantee job completion and measurable customer value. While this model in theory aligns perfectly with customer expectations for outcome-based pricing, it remains among the least popular approaches. The primary barrier is that vendors and customers struggle to agree on what constitutes a completed job or outcome. To overcome this, best practices for outcome-based models more broadly link outcomes to operational KPIs that drive business and financial performance—for example, new customers onboarded or sales closed.
- Outcome-Based: Financial Pricing.** Customers pay only for specific financial results achieved—such as revenue growth or cost reduction. This model represents the true ideal of value-aligned pricing, but it is also the hardest to operationalize. Attribution of value is challenging because vendors often do not control the full financial outcome. Consider, for instance, an AI-powered renewal agent for tech-touch customers. Even the most effective agent cannot guarantee success when a renewal depends on macroeconomic conditions, product quality, and customer satisfaction. As a result, although this model is the most disruptive, it carries the highest risks for vendors.

EXHIBIT 2
Pricing Models Are Shifting Toward Business Metrics

	 Usage-based: resources	 Agent-based	 Usage-based: interactions	 Outcome-based: jobs completed	 Outcome-based: financial pricing
Pricing based on...	Resources used or cost to serve	Number of agents deployed (proxy for user seats)	Volume of interactions or workflow steps that the agents perform	Delivery of business outcomes that support a use case	Financial performance from agentic activity
Examples	Charge based on large language model tokens, storage, compute power, or other resources	Separately monetize each type of agent (e.g., recruiter or payroll)	Monetize each interaction (e.g., conversation, ticket, execution, or API call)	Monetize activities such as leads generated, automations run, or reports created	Tie to financial outcomes, such as cost savings or increased revenue
Suited for markets where...	<ul style="list-style-type: none"> Agents support compute-heavy processes Customers value flexibility 	<ul style="list-style-type: none"> Vendors sell a large number of agents Agents perform well-defined tasks and act as labor replacement Customers value predictability 	<ul style="list-style-type: none"> Customers derive value from human-AI or agent-agent interactions such as chatbots and B2B transactions Vendors measure interaction volume 	<ul style="list-style-type: none"> Use cases and workflows are clearly defined Customers prioritize paying for what agents actually do 	<ul style="list-style-type: none"> Vendors have a high degree of control over the full outcome for the AI agents Customers are focused on clear ROI and justifying AI investments
Estimated share today and trend	40% 	20% 	25% 	10% 	<5% 

Sources: Market participant interviews; BCG analysis.

Looking ahead, we expect a continued shift away from traditional seat-based licensing toward agent-based and outcome-based pricing models. However, most incumbent software providers will not be able to transition entirely to these models without putting revenue at risk. In the near term, we anticipate broad adoption of hybrid approaches that blend established subscription structures with emerging agentic pricing models—allowing vendors to manage revenue risks while offering customers greater predictability.

Among the emerging models, agent-based pricing is likely to become more common as developers increasingly build multiple, task-specific subagents, as agents begin to substitute for human labor, and as the volume of agents deployed increases. Outcome-based financial pricing will gain traction among specific vendors, such as those positioning themselves as managed-service providers. Meanwhile, outcome-based models tied to jobs completed are likely to become increasingly common as customers focus on measuring the activities an agent completes to justify its value.

Ultimately, we expect that vendors will combine models—most commonly, agent-based and outcome-based for jobs completed. This will allow vendors to monetize both the capabilities of the agent and the value of the work it delivers.

Implementation Insights: Lessons from Early Movers

Early adopters of agentic pricing models provide valuable insights into practical implementation, highlighting both the promise and complexity of these innovative approaches.

- A provider of AI-powered sales development representatives initially adopted a usage-interaction pricing model, with charges per lead generated, email sent, and AI seat. Many users raised ROI concerns, noting that costs accumulated rapidly, often without a clear correlation to lead quality or conversion rates. Consequently, the company is now piloting a pay-for-performance model—charging per qualified opportunity—reflecting the broader industry shift toward outcome-oriented pricing.
- Similarly, a vendor of customer relationship management systems first launched its agentic offering with a fixed fee per interaction. This pricing structure soon led to customer dissatisfaction, as short, simple interactions were charged at the same rate as longer, complex ones, raising customer costs without correlating to realized value. This lack of predictability compelled sales representatives to frequently apply substantial discounts to close deals. As a result, the vendor is now transitioning to charging for discrete AI actions—such as executing a workflow or sending a reply—where users pay only when actions drive value.

- A customer service software vendor adopted an outcome-based pricing model, charging clients when an AI agent successfully resolves a customer query without human intervention. Each plan includes a baseline set of free AI resolutions, with the vendor providing tools to monitor usage and cap the number of resolutions once predefined limits are reached. Although customers generally appreciate this direct alignment of pricing to value, there have been customer disagreements about whether an issue was truly resolved. To address these challenges, the company is collaborating with customers to explicitly define what qualifies as a resolution, supplemented by an administrative arbitration process to handle disputes.

These early experiences underscore the critical importance of closely aligning pricing structures with measurable value and transparent outcome definitions to maximize customer acceptance and satisfaction.

Future-Proofing Pricing Models: Mitigating Seat Compression and Revenue Volatility

Once vendors have established an appropriate pricing structure for their agentic AI offerings, the next critical step is addressing the risk of license volume compression—in particular, the reduction in the number of user licenses or seats licensed. Rather than abruptly discarding seat-based models, vendors should embark on a gradual, multistep transition.

Initially, vendors should introduce their chosen agentic pricing structure in parallel with their existing seat-based model. This hybrid pricing model helps to diversify revenue streams and future-proof pricing strategies. Essential to this transition is developing robust tracking and telemetry capabilities that enable vendors to monetize new agentic features, provide transparent usage insights to customers, and integrate usage data into billing systems.

Determining the optimal pace of shifting from seat-based to agentic pricing requires tracking several leading indicators that signal either potential seat compression or healthy revenue diversification. Common indicators include greater agentic AI usage, purchases of fewer seat licenses, and workforce reductions within customer organizations or targeted industries. Monitoring these indicators enables vendors to make informed adjustments, progressively reducing their reliance on seat-based revenue as agentic adoption grows.

This stepwise approach ensures that vendors capture value while simultaneously laying the groundwork for scalable, transparent pricing in the future. It provides customers with the necessary

runway to build trust in agentic solutions and observe ROI—thereby reducing friction and churn while creating upselling opportunities.

To position themselves for success in this transition, vendors should focus on three priorities:

- **Define a purposeful transition.** Avoid rushing or surprising your customers, teams, and investors. Test ideas with your strategic customers, clearly communicate intentions, timelines, and the value of the new approach, and build in time for the transition in your implementation plan.
- **Build capabilities.** Develop the foundational capabilities needed to support the new pricing models—such as baselining and measuring value, usage forecasting, telemetry, quoting, and billing.
- **Evolve go-to-market models.** Because AI and agentic uptake is uneven, go-to-market models must prioritize customer success and adoption across the life cycle. For complex use cases, vendors may need to invest in resources such as forward-deployed engineering to support implementation and accelerate value realization.

These strategic initiatives collectively ensure a smooth, structured transition, mitigating the risks associated with seat compression while capitalizing on the substantial opportunities presented by agentic AI.

The rise of AI agents represents a transformative moment for the software industry, reshaping customer expectations and redefining value delivery and monetization. While adopting AI agents at scale will take time, immediate action on pricing strategies is essential in order to begin diversifying revenue and determining customer preferences. Software vendors that start adjusting their pricing model today can quickly benefit from clearer value alignment, improved customer satisfaction, and enhanced competitive advantage. Begin planning your next steps now—your strategic decisions today will directly shape your market position and profitability tomorrow.

The authors thank their BCG colleagues Arthur Libanio and Anshul Choudhari for research support.

Authors



Naila Dharani

Partner
Los Angeles



John Pineda

Managing Director & Partner
San Francisco - Bay Area



Angel Andeyro

Managing Director & Partner
New Jersey



Nate Hamming

Project Leader
Los Angeles



Yerim Kim

Principal
Los Angeles



Nipun Misra

Managing Director & Partner
Silicon Valley - Bay Area



Akash Bhatia

Managing Director & Senior
Partner, Global Sector Leader,
Technology
Silicon Valley - Bay Area



ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

© Boston Consulting Group 2025. All rights reserved.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [Facebook](#) and [X \(formerly Twitter\)](#).