CURLY'S CONSULTING / WHITE PAPER 09

Choosing the right value metric

Ammanuel "Tom" Selameab, CPP

Certified Pricing Professional

Curly's Consulting

September 2026

How to select a charging unit that follows customer value and produces workable bills.

Abstract

A charging metric determines what the customer pays for as the relationship changes. It can allow revenue to grow with useful adoption, or it can make the bill follow an activity that has little connection to customer benefit. The choice affects demand, expansion, cost recovery, and the buyer's willingness to commit.

The best metric is not necessarily the one most tightly correlated with an abstract measure of value. Buyers also need to understand, observe, budget, and influence what they are charged. The vendor needs a reliable count, an economic delivery model, and a defensible explanation for unusual cases.

For private equity owners, changing the metric can be more consequential than changing the rate. It can alter how growth translates into revenue, how automation affects the business, and how customer value is shared. Evaluate it as a redesign of the commercial relationship, including the transition from existing contracts.

The decision this paper supports

Select a charging unit by testing its relationship to buyer benefit, its effect on behavior, the distribution of bills and costs, and the feasibility of measurement. Then choose rates and commitments around that unit.

1 Separate the value driver, charging unit, and tariff

These three concepts are related but different. A value driver explains why the customer benefits, such as faster completion of an important workflow. A charging unit is what the vendor counts, such as users, processed orders, or completed cases. A tariff is the rule translating that count into expenditure.

The customer may value increased revenue while paying per transaction. The transaction is a proxy that can be observed; revenue improvement may be difficult to attribute or verify. A subscription may price access to a valuable capability even when individual uses are infrequent.

Avoid assuming that outcome pricing is always superior because it sounds closer to value. An outcome can depend on the buyer's staffing, data, product quality, or follow-up. If the vendor cannot observe those factors, the metric can create disputes or place uncontrolled risk on either party.

Likewise, a usage metric is not automatically value-based. A system that requires more clicks, retries, or API calls may generate more billable activity while providing the same result. Charging more for the vendor's own inefficiency can weaken the customer relationship.

2 Identify how customer benefit changes with scale

Map the workflow and ask what makes one customer receive more benefit than another. The relevant driver may be the number of people collaborating, the volume of work completed, the number of locations supported, the resources managed, or the importance of an occasional decision.

Then examine how the driver changes over time. Does customer growth expand the count? Does better automation reduce it? Does the product encourage consolidation into fewer accounts or administrators? A metric that appears suitable at initial sale can become misaligned as the product succeeds.

Use buyer interviews and operating data together. Historical correlations can suggest candidates, but they do not establish causal value. Large customers may have more users, more transactions, and higher revenue simultaneously. Research should identify which relationship the buyer recognizes and which measure supports a workable commercial agreement.

Consider the economic buyer's budget structure. A department may control a per-user software budget but not a share of company revenue. A customer may accept usage variability within a project while requiring a fixed annual commitment for a core operating system. Purchase feasibility belongs in the metric decision.

3 Apply gates before using a weighted score

A candidate can look attractive on average and still fail a critical requirement. Use non-negotiable gates before comparing softer tradeoffs. A count that cannot be reproduced or an outcome that cannot be attributed should not win because it scores well on perceived value alignment.

Criterion	Question to test	Evidence required
Value relationship	Does more of the unit usually mean more benefit?	Workflow and buyer evidence by segment
Comprehension	Can the buyer explain the charge?	Offer-comprehension testing
Observability	Can both parties reproduce the count?	Reconciled records and definitions
Budgetability	Can likely and high bills be estimated?	Account-level bill distributions
Control and incentives	Does the charge discourage useful behavior?	Interviews and observed response
Delivery economics	Does revenue cover plausible cost variation?	Contribution under normal and heavy use
Commercial fit	Can the organization buy under this structure?	Budget, procurement, and contract evidence

Where several candidates pass the gates, compare them transparently. Weight criteria according to the strategy and sales motion, but preserve the underlying evidence. A score of four out of five is a judgment summary, not an empirical result or a substitute for testing the bill.

4 Evaluate metric and tariff together

A per-unit metric can support a flat usage rate, volume tiers, an annual commitment, or a base fee with an allowance. These offers expose the buyer and vendor to different risks even when the same unit is counted.

Two-part tariffs combine an access fee and a usage price, a structure explored in Oi's foundational analysis.[1] A three-part tariff adds an included allowance before overage charges begin. The combination can support cost recovery and bill predictability, but the commercial terms still require evidence.

Lambrecht and colleagues' study of three-part tariffs finds that uncertainty about usage affects plan choice in its setting.[2] Ascarza and colleagues show that tariff structure can also affect usage beyond a simple mechanical budget effect.[3] These findings support modeling behavior under the proposed tariff rather than assuming that historical usage will remain fixed. They do not establish a universal response for B2B products.

For a bill function $B(q)$ and delivery cost $C(q)$, expected contribution depends on the distribution of usage under the offer:

$$\mathbb{E}[\text{Contribution}] = \mathbb{E}[B(q) - C(q)] - S,$$

where S is incremental account-level operating cost not already in C. If the tariff changes usage, use the new distribution rather than merely applying a new rate to old activity.

Inspect low-usage economics, the median, the high-usage tail, and customers near thresholds. Include failed work and retries in cost even if they are not billable. A metric that tracks accepted output can align with customer value while requiring the vendor to absorb the cost of unsuccessful attempts.

5 Worked example: automation changes the seat count

Illustrative assumptions. A product supports 10,000 completed tasks each month. The customer initially has 50 relevant users, and later has 20 after workflow automation while completed volume rises to 15,000 tasks. These are hypothetical operating states, not a causal forecast of staffing effects.

Compare three candidate offers: $40 per user per month; $0.12 per completed task; and a hybrid of $800 per month plus $0.08 per completed task. Assume delivery cost averages $0.03 per completed task, inclusive of associated failed attempts, and exclude shared and account-level operating expense for simplicity.

Candidate offer	Initial bill	Later bill	Initial contribution	Later contribution
Per user	$2,000	$800	$1,700	$350
Per completed task	$1,200	$1,800	$900	$1,350
Hybrid	$1,600	$2,000	$1,300	$1,550

The seat metric loses revenue as user count falls, even though work completed increases. The task and hybrid offers behave differently. That does not prove either is the right choice. At the stated rates, both reduce initial contribution compared with seats, and the buyer may perceive a variable task bill differently from a predictable seat bill.

The research questions are concrete: Do completed tasks reflect value in this workflow? Can the buyer verify and budget the count? Does a per-task charge discourage useful work? What rates and commitments will the buyer accept? Will actual behavior differ once usage is charged?

The example isolates structural tradeoffs. It deliberately does not assume that changing the metric preserves demand, retention, usage, or negotiated price. Those responses must enter the commercial case before migration.

6 Define the billable event precisely

For every usage or outcome metric, write a specification that the customer and vendor can apply to the same records. Include the event, eligibility, timing, aggregation level, exclusions, and correction policy.

For a resolved support case, define what constitutes resolution, whether customer confirmation is required, how reopened cases are treated, and whether human-assisted outcomes count. For an active user, define activity, unique identity, administrators, guests, and duplicate accounts. For processed transactions, define failures, reversals, tests, and refunds.

Decide who controls the event. If the vendor can increase billable events through retries or splitting work, the count may create a conflict. If the buyer can avoid the count while retaining the benefit, the model may be easy to circumvent. Address these incentives through the definition and product design rather than relying only on enforcement.

Build a dispute path and an auditable usage record. The buyer should be able to understand the bill without access to proprietary internal infrastructure. A comprehensible account-level activity statement is often more valuable than an opaque count presented with absolute confidence.

7 Research total bills, not isolated unit prices

A buyer reacting to "twelve cents per task" may not recognize the annual expenditure implied by its workload. Present realistic low, expected, and high usage scenarios alongside the unit rate, commitments, and overages. Confirm that respondents understand the bill before interpreting their preference.

Compare credible alternatives with equivalent scope. A flat subscription and a usage offer should be presented as complete commercial propositions, not one detailed offer against an underspecified alternative. Include the current contract where incumbents are being studied.

Investigate the value of predictability. Some buyers will prefer a higher expected payment with a credible ceiling or a committed allowance. Others will prefer a variable bill that falls when demand falls. The preferred balance can differ by segment and budget process.

Research the user's behavior separately from procurement approval. A unit that a budget owner accepts may discourage daily adoption if users perceive every action as an expense. Conversely, users may generate substantial variable cost that the buyer only discovers at invoice time. Both perspectives matter.

8 Plan migration with shadow bills and behavioral validation

Use historical usage to produce shadow bills under candidate structures. Compare the distribution of changes across accounts, including customers with seasonal demand and unusually heavy use. This exercise reveals mechanical exposure and helps identify where transition terms may be needed.

Shadow billing does not estimate the behavioral response to the new charge. Historical users did not face the candidate tariff. Follow with a commercial test or staged introduction where feasible, observing adoption, spend, disputes, and retention. Preserve a credible comparison when assessing the effect.

Map contractual eligibility and the customer effort required to move. A new count may require instrumentation, reporting access, purchase-order changes, or internal budget approval. Include these costs and timing in the investment case.

Decide whether new customers and incumbents should transition on different schedules. Temporary price protection or commitments may improve adoption but delay financial benefits. Make those terms explicit and give them an owner and expiry rather than allowing a permanent collection of bespoke arrangements.

9 Choose for the commercial relationship you want to build

The final recommendation should explain why the metric fits the buyer's value, how the bill behaves, where contribution is exposed, and what evidence supports adoption. It should also identify the scenarios in which the metric becomes less suitable as the product changes.

Review the metric when automation, workflow scope, customer scale, or delivery costs change materially. Do not assume that a successful historical charging unit remains suitable indefinitely. At the same time, avoid frequent redesign without evidence; each change imposes learning and migration costs.

A useful value metric makes successful use commercially sustainable. It lets the customer understand the exchange, gives the supplier an appropriate way to participate in growth, and keeps the cost of serving that growth visible. The right choice emerges from that complete relationship, not from a fashionable pricing label.

10 Research the unit customers will accept and the business can sustain

The missing evidence is usually a combination of buyer and operating data: how value varies, what the buyer can budget and control, how usage is distributed, what different bills imply, and how behavior changes when the unit is charged. A competitor's metric does not establish those relationships for your product.

Commission a comparison of complete commercial offers. Require metric-comprehension testing, buyer tradeoffs, account-level bill and contribution simulations, and a migration analysis. Where behavior remains uncertain, specify the commercial test that will resolve it. The result should connect the unit, tariff, customer response, and economics rather than recommend a metric in isolation.

Curly's Consulting's **Pricing Transformation** combines primary buyer research, willingness-to-pay analysis, monetization design, size-of-prize modeling, and support with execution. For a metric decision, the relevant scope should make the value relationship and bill distribution visible before the company changes its contracts.

The next commercial step

Request a Pricing Transformation quote. Include the current metric, alternatives being considered, customer workflow, and usage or cost data already available.

Sources and methodology

This paper combines the cited published work with Curly's Consulting's analytical framework. The commercial examples, thresholds, and calculations are illustrative; they are not client results or findings from a new empirical study. Published studies support the specific claims beside their citations. Their results should not be treated as universal benchmarks for a portfolio company.

References

[1] Oi, W. Y. (1971). A Disneyland dilemma: two-part tariffs for a Mickey Mouse monopoly. *Quarterly Journal of Economics*, 85(1), 77–96. https://doi.org/10.2307/1881841.

[2] Lambrecht, A., Seim, K., and Skiera, B. (2007). Does uncertainty matter? Consumer behavior under three-part tariffs. *Marketing Science*, 26(5), 698–710. https://doi.org/10.1287/mksc.1070.0283.

[3] Ascarza, E., Lambrecht, A., and Vilcassim, N. (2012). When talk is "free": the effect of tariff structure on usage under two- and three-part tariffs. *Journal of Marketing Research*, 49(6), 882–889. https://doi.org/10.1509/jmr.10.0444.

About the author

Ammanuel "Tom" Selameab, CPP, is a Certified Pricing Professional and the founder of Curly's Consulting. He helps investors and operators make pricing and monetization decisions through primary buyer research, financial modeling, and commercial analysis.
tom@curlysconsulting.com | curlysconsulting.com

© 2026 Curly's Consulting. All rights reserved.