A Practical Guide to Adopting Usage-Based Pricing Without Breaking What Works

With insights from
James D. Wilton
CEO and Senior Partner,
Monevate

What You'll Find Inside

Pricing is the New Product Narrative 02

Usage-Based Pricing's Second Act: Why It's Having Another Moment in the Spotlight 03

When Should You Consider Usage-Based Pricing? 04

The Million-Dollar Question: What Should You Meter? 05

Have Your Cake and Eat It Too: Layering Usage Into Your Pricing Strategy 06

Designing Customer Experience That Builds Trust and Protects Revenue 07

The People Side of Pricing: Who Does What When You Go Usage-Based 08

Infrastructure That Makes Usage-Based Pricing Work at Scale 09

Moving Customers to New Pricing Without Losing Them 10

Proof of Concept: How to Know If Your New Pricing Is Working 11

Pricing is the New Product Narrative

Creating value is easier than ever. AI-native features roll out in days, workflows launch without code, and cloud infrastructure scales instantly.

Shipping something useful is no longer the hard part. **Capturing value is.**

80% of companies still take one quarter or more to test pricing or even align on the right value metric.

By then, the feature's out, the value's been delivered. And by the time pricing catches up, the moment to capture that value has slipped away.

The gap between value created and value captured is widening, and nowhere is it more exposed than in pricing. The old models aren't holding up. Seats don't map to value. Bundles force customers to pay for features they never use. And buyers increasingly expect pricing that reflects outcomes, not just access.

Usage-based pricing sounds like a model shift. But in practice, it's a company-wide transformation that touches product, sales, customer success, finance, and even culture. That's what this playbook is about.
It's not a pitch for usage-based pricing. It's a tactical guide to make it work: how to structure your pricing, what infrastructure you'll need, and how to migrate without losing customers' trust.

Ready to make pricing a key driver of your growth? This is where you begin.

So you start thinking about usage-based pricing. But then the real questions hit:

- What's the right value metric?
- How do we comp sales if the revenue trickles in post-close?
- What happens to legacy customers on fixed plans?
- Can our product even track usage at the right level?

1

Usage-Based Pricing's Second Act: Why It's Having Another Moment in the Spotlight

Usage-based pricing (UBP) is not a new concept, but its relevance has skyrocketed in recent years, thanks to the rise of AI. It charges customers based on their actual consumption, rather than just access. UBP can take the form of pure pay-as-you-go (e.g., per API call, GB used, or images created) or be blended with fixed tiers or subscriptions.

The Pre-AI Momentum

Before AI dominated headlines, SaaS companies with usage-based components already saw advantages: higher net revenue retention as pricing scaled with customer growth; lower barriers to entry for hesitant prospects; and more organic expansion as usage increased.

Customers appreciated the transparency. Pricing that reflected actual consumption felt more defensible than arbitrary seat counts or feature bundles.

AI's Push for Value-Based Pricing

Despite the rise of usage-based pricing, most software products had a fixed subscription model: an upfront platform fee and user seats as the primary value metric. But the way AI delivers value (highly variable, output-driven, and scaled by user intent) has forced companies to pause and reassess how they monetize their products.

1. **Seat-Based Models Under Pressure**

   James notes, "**It used to be the case that the amount of value I got from the product was very correlated with the number of people I had working in it, which is why per-user pricing worked. But now, with AI and automation, that doesn't hold anymore.**"

   A single user with AI assistance can accomplish a lot more work. Automated workflows handle tasks without human involvement. AI-native companies are scaling with leaner teams than traditional software companies ever did.

   The result? Revenue tied purely to headcount struggles to capture the expanding value these companies create. It's not that seats are completely irrelevant; they're now insufficient as the primary value metric for most companies.

2. **Variable Value, Fixed Pricing**

   AI capabilities deliver wildly different value depending on the complexity of the use case, the depth of integration, and the sophistication of the model. One customer might achieve a 10x ROI, while another sees only marginal benefit.

   Fixed pricing treats both identically, creating misalignment. High-value users may feel they're getting a bargain while low-usage customers question the cost.

3. **Every AI Feature Behaves Like a Microproduct**

   AI features consume real, variable costs: tokens, compute, storage, bandwidth. These costs fluctuate based on input complexity, output type (text/audio/video), and usage patterns.

   Treating resource-intensive AI capabilities as flat monthly features creates unsustainable unit economics as usage scales.

Key takeaway

Whether you're selling to lean AI-first startups or enterprise teams embracing automation, one thing is clear: Your pricing model should reflect how your product is used and where value is delivered, not just the number of people who have access to it.

2

When Should You Consider Usage-Based Pricing?

As with any pricing model, usage-based pricing isn't a silver bullet or universal fix for monetization. However, it's essential to regularly assess whether your current pricing model continues to meet your needs.

James outlines a core truth: "If you're a fast-growing SaaS company, your pricing strategy will probably only last you 2–3 years, because your product, your market, or your goals will change. With AI, that shelf life compresses even further."

When that shift comes, don't just tweak price points. Start with one key question: Does your pricing reflect how customers experience value today?

If the answer isn't clear, use this checklist.

The UBP Fit Check

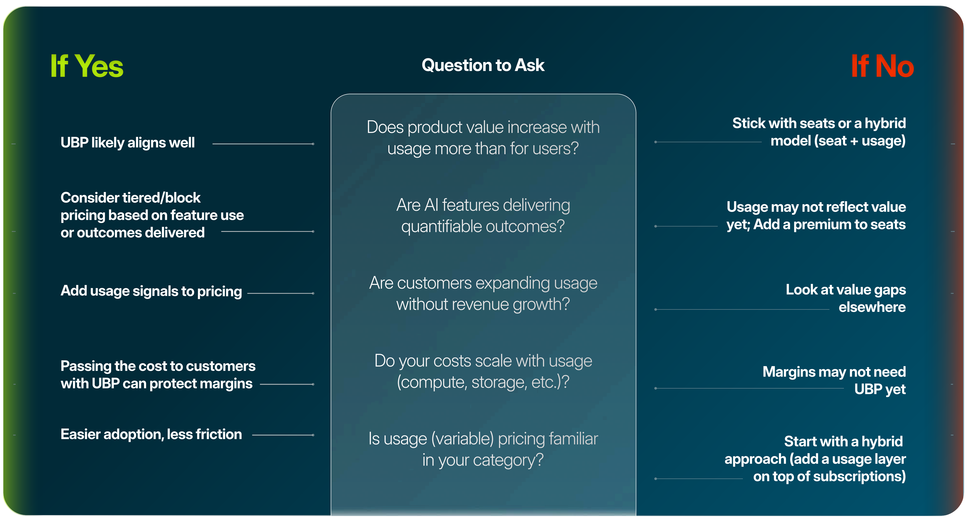

If multiple signals apply, UBP may be more than just viable; it might be necessary. Don't wait for revenue to plateau or costs to balloon. Evaluate early, and experiment with intent.

> ***The general idea of moving to usage-based pricing for AI is... (surprise, surprise!) to map the value that our customers are getting from AI capabilities to pricing, in a way that's more tightly coupled than what a seat-based model enables. It's enabling customers to pay for what they actually need while enabling the company to offset the costs of serving AI models and maintain healthy margins.***
>
> **Tony Beltramelli,**
> Head of Product for AI,
> **Miro**

Plot Twist: What if Your Customers Are Already Living Usage-First?

Monte Carlo, a data and AI observability platform, didn't have to convince customers to accept usage-based pricing—they had to stop frustrating them with upfront annual contracts. Their buyers were developers used to AWS, Datadog, and Snowflake. Paying only for what you use wasn't novel; it was expected. So Monte Carlo pivoted to a pure pay-as-you-go model and started tracking 'daily revenue'. The shift landed instantly and became core to their GTM strategy.

Key takeaway

Usage-based pricing works when your product's value scales with usage, your customers understand that link, and your teams are ready to sell, support, and charge for it.

3 The Million-Dollar Question: What Should You Meter?

One of the most critical decisions in usage-based pricing is deciding your value metric(s) (what to meter). Your value metric becomes the foundation for customers' understanding of their bills and whether they believe the pricing is fair.

Choose poorly, and customers feel confused or exploited. Choose well, and pricing feels natural and defensible.

Start with Value Assessment

Before selecting metrics, document where your product creates specific value for customers. Then, through customer interviews and surveys, test which value dimensions matter most to different segments.

Four Types of Value Metrics to Consider

- Inputs: Data volume, user queries, API calls, tokens consumed
- Usage: Sessions, GPU hours, processing time
- Outputs: Documents generated, words translated, workflows completed
- Outcomes: Time saved, costs avoided, new revenue unlocked

Evaluate Your Options: What Makes a Good Value Metric?

Once you've identified potential metrics across these categories, put each candidate through this evaluation framework. The metrics that score well across most criteria are your strongest contenders:

Criteria	Question to Ask	Why It Matters
Value Alignment	Does the metric reflect how customers get value?	If pricing feels disconnected from outcomes, customers will resist.
Growth Orientation	Will this metric scale with customer success?	The best metrics grow in tandem with customer adoption and usage.
Clarity	Is it easy for customers to understand and predict before they buy?	Complex metrics create friction in sales cycles.
Auditability	Can both sides verify usage post-purchase?	Billing disputes over usage charges erode trust.
Predictability	Can customers forecast their costs with confidence?	Unpredictable bills create anxiety and churn risk.
Controllability	Can users adjust their spending?	Giving users control over their spend prevents fear of runaway costs.

Avoiding the AI 'Black Box' Trap:

With AI products, customers resist paying for processes they can't see or understand, especially when those processes appear to be system inefficiencies rather than value delivery.

The root cause of this mismatch is that companies often price based on what drives their cost, not what drives customer value. "In AI, the instinct is to charge per query or token," James says. "That helps recover cost, but customers don't think in tokens. They refine queries, backtrack, and iterate. Value lives in the final output, not every input."

He recalls a search provider that charged per query, but every keystroke triggered a new one. Typing a two-word search could result in a dozen events before hitting enter, and customers felt cheated.

The fix? Charge per session instead. One billable unit for the full intent. This addressed customers' perception of being nickel-and-dimed.

Choose metrics that are visible to customers and align with their perception of value. Charge for documents processed rather than tokens consumed, or results delivered rather than queries executed. Focus on what makes your product or customer results distinct, rather than just general AI processing capabilities.

When Zapier Bet on the Wrong Value Metric (And How They Fixed It)

Zapier once priced based on Zaps (automations) and Tasks (each run), but users often hit the limit on one metric while still having quota left on the other, forcing them to upgrade without feeling they'd fully used their plan. This disconnect led to frustration, confusion, and mounting complaints. Zapier overhauled its pricing to prioritize trust, removing arbitrary limits, adding pay-as-you-go flexibility, and simplifying the model. A year later, usage surged, churn dropped, and metered revenue grew.

Read the full story →

Key takeaway

Your value metric shapes the entire customer relationship. It determines whether pricing feels fair, whether customers can predict costs, and whether growth feels sustainable for both parties. Take time to align your metric with customer value, not just your cost structure.

4 Have Your Cake and Eat It Too: Layering Usage Into Your Pricing Strategy

Choosing what to meter is just the beginning. The more complex question is how to introduce usage without abandoning your stable revenue streams

This is where many companies get stuck. When usage-based pricing enters the conversation, it tends to feel like a fork in the road: stick with your existing subscriptions or swing hard into pure pay-as-you-go (which often sounds 'highly volatile').

In reality, the best approach is usually a hybrid one: adding usage-based components on top of your predictable subscription elements.

And the data backs this up. According to our 2025 State of Subscriptions and Revenue Growth Report, **43% of companies now combine subscriptions with usage-based pricing**. Pure-play models still exist, but the trend is unmistakable: most businesses are layering usage, not replacing subscriptions.

Why Hybrid Models Dominate

Two major forces drive companies toward hybrid pricing:

1. **Buyer psychology:** Even usage-friendly customers want predictability for budgeting and procurement. As James notes, "Enterprise buyers are often willing to pay **15% more** just to have certainty around what they'll owe. That level of predictability gives them confidence to move forward."

2. **Organizational reality:** James explains, "Selling based on usage is completely different from selling based on users. Your salespeople need to speak a different language. Your CS team needs to drive usage. Your whole organization needs to pivot, so doing it all at once is rare."

4 Popular Hybrid Models (And When to Use Them)

1. **Base plan + Overage:** Customers pay a flat fee for a defined usage quota, then pay per unit for overages (often at premium rates).

   Example: *Zapier includes set tasks in each plan, charging 1.25 times the base rate for extras.*

   - **Buyer perspective:** "I know my minimum cost and only pay more if I grow."
   - **When to use:** Ideal when you want a predictable entry price with room to scale. Works best when usage is easy to explain and overages are infrequent but expected.

2. **Block-Based Capacity:** Customers pre-purchase usage blocks rather than paying per unit in real-time.

   Example: *Phrase offers clear pricing plans with included capacity (e.g., the number of words stored or translated), allowing customers to purchase additional blocks as needed.*

   - **Buyer perspective:** "I don't have to worry about a meter running. I scale in chunks when I'm ready."
   - **When to use:** Best for PLG or sales-led motions where customers value predictability but need flexibility to grow without committing to a full-tier upgrade.

3. **Minimum Commitment + True-Up:** Customers commit to a minimum spending amount for discounted rates, with true-up charges applied if they fall short of the commitment.

   Example: *Common in platforms like AWS or Splunk, a customer commits to $20,000 for Q1. By quarter's end, if they've only used $18,000, a $2,000 true-up is invoiced.*

   - **Buyer perspective:** "I get better rates for promising to spend, and I know exactly what I'll owe if I don't use it fully."
   - **When to use:** Ideal for enterprise or high-usage accounts that require locking in revenue while offering flexible usage. Useful when the buyer's procurement needs a clear commitment number.

4. **Prepaid Credits & Drawdown:** Customers pre-purchase a pool of credits tied to a monetary value, which they draw down as they consume services. You define how credits convert to usage, which features or limits they unlock, and whether unused credits roll over to the next billing period.

   Example: *Freepik allows users to buy credit packs (e.g., 100 credits) and spend them on downloads, premium assets, or AI-generated content. Each asset has a credit cost, and the balance depletes as users consume it.*

   - **Buyer POV:** "I want to commit upfront to get better rates, but I need flexibility to ramp usage when the time is right."
   - **When to use:** Ideal for generative AI and agentic AI use cases. Prepaid credits let you monetize multiple features or agents that may have different underlying costs and values, while keeping pricing simple for users. Because all features draw from a shared pool of credits, you can offer flexibility without exposing customers to pricing complexity.

Key takeaway

Your pricing model sets the foundation for how you monetize growth. Whether you choose overages, blocks, prepayments, or hybrid commits, the goal is the same: align price with value delivered. Hybrid models give you flexibility without forcing a binary choice. Start with what's easy to explain, easy to sell, and easy to scale. Evolve as your product and customers mature.

5 Designing Customer Experience That Builds Trust and Protects Revenue

Your pricing strategy is only half the equation. The other half is how customers feel about paying you. Usage-based pricing doesn't reside in spreadsheets—it appears in invoices, upgrade prompts, alert emails, and product limits. Every billing cycle, usage threshold, and overage sends a signal: Is your pricing fair or predatory? Predictable or chaotic? Does it build trust, or erode it?

Beyond perception, billing mechanics affect cash flow, revenue recognition, fraud risk, and operational complexity.

This chapter demonstrates how to create pricing that feels fair, predictable, and scalable, without compromising financial control or customer trust.

Decide Your Billing Cadence

Once you've locked in your pricing model, the next decision is billing cadence: how often customers pay and when usage gets invoiced. Your billing rhythm has a direct impact on cash flow, risk exposure, and customer experience.

Below are the most common billing cycles and their impact on your operations.

Billing Frequency	Pros	Cons
Monthly Billing	Customer-friendly; tighter feedback loop for your teams	Less predictability; More invoices to deal with every month
Quarterly Billing	Predictable for finance; fewer invoices	Slower to detect usage spikes or react to growth patterns
Annual Contract + Monthly Usage Billing	Locks in commitment while spreading out usage charges	Requires clear thresholds and real-time usage tracking
Usage in Arrears (End-of-Term Billing)	Simple setup; good for early adoption or low-touch tiers	Higher exposure to non-payment or fraud if usage accrues without interim billing
Threshold-Based Billing	Protects against overuse without payment	Needs more sophisticated tracking and usage alerts

Pro Tip: For high-usage or free-tier customers, set a billing threshold (e.g., 50% above quota) to trigger mid-cycle invoices. This reduces fraud exposure and maintains a steady cash flow

Plan for What Happens Beyond Usage Limits

Once usage exceeds a plan's quota, how you respond matters. Some customers expect uninterrupted service and are willing to pay for it. Others prefer cost controls, even if it means being throttled or cut off. The right enforcement model depends on your product's role and your buyer's mindset.

Limit Handling Strategy	What It Does	How It Helps	Best Fit For
Soft Overages	Let users exceed limits and charge later	Keeps experience smooth; revenue keeps flowing	Mission-critical tools; trusted enterprise customers
Hard Limits (Cutoff)	Block access when the quota is hit	Protects margins; risky UX unless clearly communicated	Freemium plans; non-critical usage; resource-intensive products
Throttling	Slow down performance instead of cutting access	Maintains access while signaling the need to upgrade	Freemium or PLG products; cost-sensitive user bases
Prepaid Top-Ups	Let users buy additional usage as needed	Gives control; avoids surprise bills	SMBs or international markets with fixed budgets
Auto-Upgrades	Move users to the next tier automatically	Minimizes friction, but must be transparent	PLG or mid-market tools with clearly tiered plans

Whatever path you choose, make it transparent. Utilize alerts, dashboards, and billing previews to prevent surprises. Clearly document your usage and billing policies to establish expectations and minimize the support burden. Stay close to your customers to understand their perspective and adjust as needed.

Key takeaway

The moment a customer hits their limit or receives their bill is when pricing becomes personal. Design these moments with care. Your billing cadence, enforcement rules, and transparency don't just affect revenue; they shape trust. And trust is what usage-based models are ultimately built on. And behind every smooth experience is a system built to handle it. We'll get to that later.

6 The People Side of Pricing: Who Does What When You Go Usage-Based

Once you've nailed your pricing model and mapped out the billing mechanics, the strategy work is largely done. What comes next is execution, and that depends on people. Usage-based pricing may initially be a monetization decision, but sustaining it requires ongoing coordination across multiple teams.

Because in a usage-based world, revenue doesn't just show up at contract signature. It becomes apparent when the product is used, meaning that Product, Sales, CS, Finance, and RevOps all have a role to play. This chapter covers how to assign ownership, structure incentives, and operationalize pricing to prevent it from stalling out post-launch.

Decide Who Owns Pricing

Before you rethink roles, decide who's responsible. Pricing often starts as a founder-led decision or something debated ad hoc between product and finance. But if no one owns it, it doesn't move. You need someone obsessing over usage metrics, value alignment, and packaging experiments —someone who treats pricing as a product in itself. Whether that's a monetization lead, product ops, or a cross-functional working group, pick an owner.

Product's Foundational Role

The product is the ground zero for usage-based pricing. Product teams enable the model by:

- **Instrumenting the product** to track usage at a granular level
- **Surfacing insights** to customers so they can self-optimize
- **Prioritizing features and workflows** that unlock value tied to pricing

In usage-based models, the "value metric" isn't theoretical. It's visible in-product. And that makes product teams critical to making pricing feel fair, scalable, and outcome-aligned.

Sales Compensation Needs to Shift

Traditional sales comp plans reward upfront deal size. But in usage-based pricing, the deal only starts delivering value and revenue once the customer actually uses the product. That's a different rhythm.

- **Reward land + expand:** Comp reps on customer growth over time, not just initial contract value
- **Discourage overcommits:** Penalize inflated usage estimates that lead to refunds or downsells
- **Align with usage milestones:** Incentivize product adoption, not just deal closure

Reps don't need to become consultants, but they do need to understand usage behavior and value realization to sell confidently. A sales comp structure tied to growth also becomes a training ground, as it builds the habits and instincts needed to drive long-term value, not just quick wins.

Customer Success Becomes a Growth Driver

CS can no longer be a reactive "churn prevention" team. Their job now is to:

- **Drive usage** by helping customers discover use cases tied to value
- **Monitor telemetry** to spot drop-offs or expansion signals
- **Trigger interventions** at usage thresholds, upgrade points, or onboarding gaps

Finance Must Relearn Forecasting

With variable usage comes variable revenue. Finance needs to rethink:

- **Forecasting:** Model it based on usage signals, not fixed ACV
- **Risk controls:** Use true-ups, prepayments, or usage caps
- **Revenue recognition:** Align billing systems to handle mid-cycle charges and fluctuating usage

RevOps as the Pricing Backbone

In usage-based pricing, RevOps isn't just operational support; it's central to how pricing works at scale, connecting strategy and systems. They:

- Ensure usage data flows cleanly across CRM, billing, and analytics
- Create dashboards for tracking quota, usage, and expansion potential
- Own the logic for entitlements, thresholds, and billing triggers

Key takeaway

Successful usage-based pricing depends on how well your teams adapt to it. Sales needs new comp plans, CS needs to drive usage, Product must build metering and visibility, and Finance has to model for variability. Above all, someone needs to own the strategy and continually evolve it.

7 Infrastructure That Makes Usage-Based Pricing Work at Scale

So you have your pricing model in place, your teams aligned, and the mechanics designed. But none of it works without the infrastructure to support it. Unlike seat-based or flat subscriptions, usage-based pricing relies on real-time signals, dynamic entitlements, and precise billing, because every spike, overage, or delay is something your customer feels.

Many teams underestimate the scope. This isn't just a back-office upgrade—it's a system-wide shift across product, billing, data, and finance.

Our recent market study of 450+ SaaS and AI companies revealed a clear pattern: apart from value articulation, the biggest hurdles to adoption are all infrastructure-related.

1. Usage Ingestion – Capture Every Signal at Scale

If usage is your revenue trigger, missing or duplicated data can lead to revenue leakage and erosion of trust.

You need an ingestion system that can:

- Ingest high-volume usage data in real time or batch
- Handle diverse event types (API calls, compute, tokens, messages)
- Maintain clean, reliable data with built-in de-duplication and idempotency
- Scale horizontally to support AI or data-intensive workloads
- Allow for flexible schema changes without breaking downstream processes

Why it matters: Accurate, high-throughput ingestion is the foundation of usage-based pricing. If the data is off, everything else downstream, from metering to billing, is compromised.

2. Metering & Rating – Convert Usage Events Into Revenue

Raw events aren't billable until they're transformed into pricing logic.

Your system should:

- Let you define and iterate metered features (e.g., docs processed, GB stored) without engineering heavy lifting
- Filter and aggregate usage data using time windows and logic (like SUM or COUNT)
- Support flexible rating models (tiers, packages, thresholds) and price updates without code rewrites

Why it matters: Shifting to usage-based pricing involves discovering what customers truly value and how they're willing to pay. That takes iteration. The faster you can define and refine meters and pricing logic, the quicker you'll land on a model that works.

3. Billing Engine – Send Invoices That Inspire Trust

Usage-based pricing strains traditional billing systems. You need an engine that can:

- Calculate usage-based charges in real time or at billing intervals
- Support hybrid models (base + overage, tiered usage, drawdowns)
- Manage proration, mid-cycle changes, and multi-entity invoicing
- Integrate natively with quoting, checkout, and contracts
- Automate tax handling and support multi-currency billing
- Provide detailed audit trails and invoice previews for transparency

Why it matters: With UBP, your billing engine becomes an extension of your product. If it can't support your pricing models, your customer experience, and revenue will take a hit.

4. Entitlements – Control Access and Track Consumption

With usage-based pricing, provisioning becomes a revenue lever. You need to manage:

- Feature-level access and usage thresholds (e.g., 1M API calls/month)
- Real-time enforcement (throttle, notify, block)
- Sync between entitlements, usage data, and billing systems
- Support for upgrades, trial overrides, and plan changes
- Visibility into "purchased vs. consumed" usage

Why it matters: Entitlements connect billing to the product. They prevent overconsumption, reduce revenue leakage, and give you the flexibility to iterate on feature packaging. Learn how Phrase leveraged entitlements in its pivot to usage-based pricing.

5. Data Architecture – Align Product, GTM, and Finance

Usage data can't live in silos. You need:

- A unified source of truth for usage and billing data
- Reports and analytics to visualize usage across cohorts
- Integrations across your tech stack—CRM, CPQ, revenue recognition, and accounting tools

Why it matters: When usage and billing data are disconnected, teams lose alignment and confidence. Choose a system that brings them together so everyone sees the same numbers and knows they can trust them.

Key takeaway

Your monetization strategy is only as agile as your infrastructure. While some tools capture usage data and others handle billing, fragmented systems slow down pricing experiments. Tools like Chargebee unify ingestion, metering, entitlements, and billing, so you can iterate quickly without patching together point solutions.

8

Moving Customers to New Pricing Without Losing Them

Changing pricing for new customers is more straightforward because they don't have baggage. Existing customers do. Contracts, expectations, and a mental model shaped by your old pricing. That's what makes migration tricky. This isn't a flip-the-switch moment. It's a structured rollout to realign value and pricing without breaking trust.

James notes, "If you plot what a customer is paying today vs. what they'd pay under your new usage model, it usually looks like someone fired a shotgun at the graph—completely scattered. And it's natural because your old value metrics, like seats, rarely map cleanly to usage."

That's why successful migrations depend on intelligent segmentation and pricing guardrails.

Start With New Customers

Rolling out your new pricing model with net-new customers lets you learn without the risk. There are no expectations to reset, and every closed deal builds internal muscle, providing your team with real data, tested talk tracks, and the confidence to engage with legacy customers later.

Set Pricing Boundaries: Target vs. Floor

- Target price is what a customer would pay under the new model based on actual usage.
- Floor pricing is your minimum acceptable threshold: what you're willing to offer during migration, while protecting margin and pricing integrity.

These two numbers create room to negotiate without compromising strategy.

Phase the Rollout

Think of migration as a campaign, not a cutover. Use a phased approach to reduce risk and build momentum:

Phase	Segment	How to Identify Them	Focus
1	New customers (No baggage)	-	Launch the new pricing model and learn
2	Low-risk legacy accounts	Small deals, strong relationships, or upcoming renewals	Test messaging and build processes for broader rollout
3	Mid-tier accounts	Moderate complexity and usage patterns	Scale migration with learnings from earlier phases
4	High-value customers	Large ARR, bespoke contracts, executive involvement	Use custom rollout plans with tight exec alignment
5	Remaining base	Legacy accounts with edge-case usage or no urgency	Standardize and automate where possible

Key takeaway

Transitioning your base is less about enforcing a new price and more about resetting the value conversation. Do it gradually, with structure, and you will earn the trust of your customers.

9

Proof of Concept: How to Know If Your New Pricing Is Working

Once your new pricing is live, the real question is: is it working? Not just “are we billing correctly,” but “are we seeing the results we designed for?”

This chapter provides the metrics and signals to track, aligned with the goals that likely triggered your initial shift.

Goal	Key Metrics
Expand revenue	• Net Revenue Retention (NRR) • Average Revenue Per Account (ARPA) • Upgrade velocity (Time from initial purchase to first expansion)
Shorten sales cycles	• Days to close • Conversion rates at each stage
Improve predictability	• Number of billing disputes/mo. • Forecast accuracy

Key takeaway

You won’t get pricing perfect on day one, and that’s fine. What matters is staying close to how customers experience value, tracking what moves the needle, and adjusting as you learn. Progress in usage-based pricing comes from iteration, not certainty.

About the Contributors

James D. Wilton is the CEO and Senior Partner of Monevate, a strategic advisory firm focused on monetization, pricing, and business model design for high-growth SaaS and AI companies. He has helped scale pricing strategies at some of the most innovative startups and growth-stage businesses, working directly with product, finance, and GTM leaders to bridge the gap between value creation and value capture.

James is recognized for bringing clarity to complex pricing transitions, especially in usage-based and hybrid models, and for helping teams to move faster without compromising customer trust.

Harikrishna is a Senior Product & Solutions Marketer at Chargebee, the leading Revenue Growth Management platform that helps over 6,500 companies manage multiple pricing models, launch freemium offerings, move upmarket and sell to enterprises, invoice and recognize revenue accurately, and prevent subscriber churn. He works at the intersection of product strategy, customer insight, and go-to-market execution, translating complex problems into actionable narratives that resonate.

At Chargebee, Harikrishna leads key initiatives around usage-based pricing and AI monetization, partnering closely with SaaS and AI innovators to navigate pricing transformation. This playbook reflects his ongoing work to help companies build scalable and story-worthy monetization strategies.

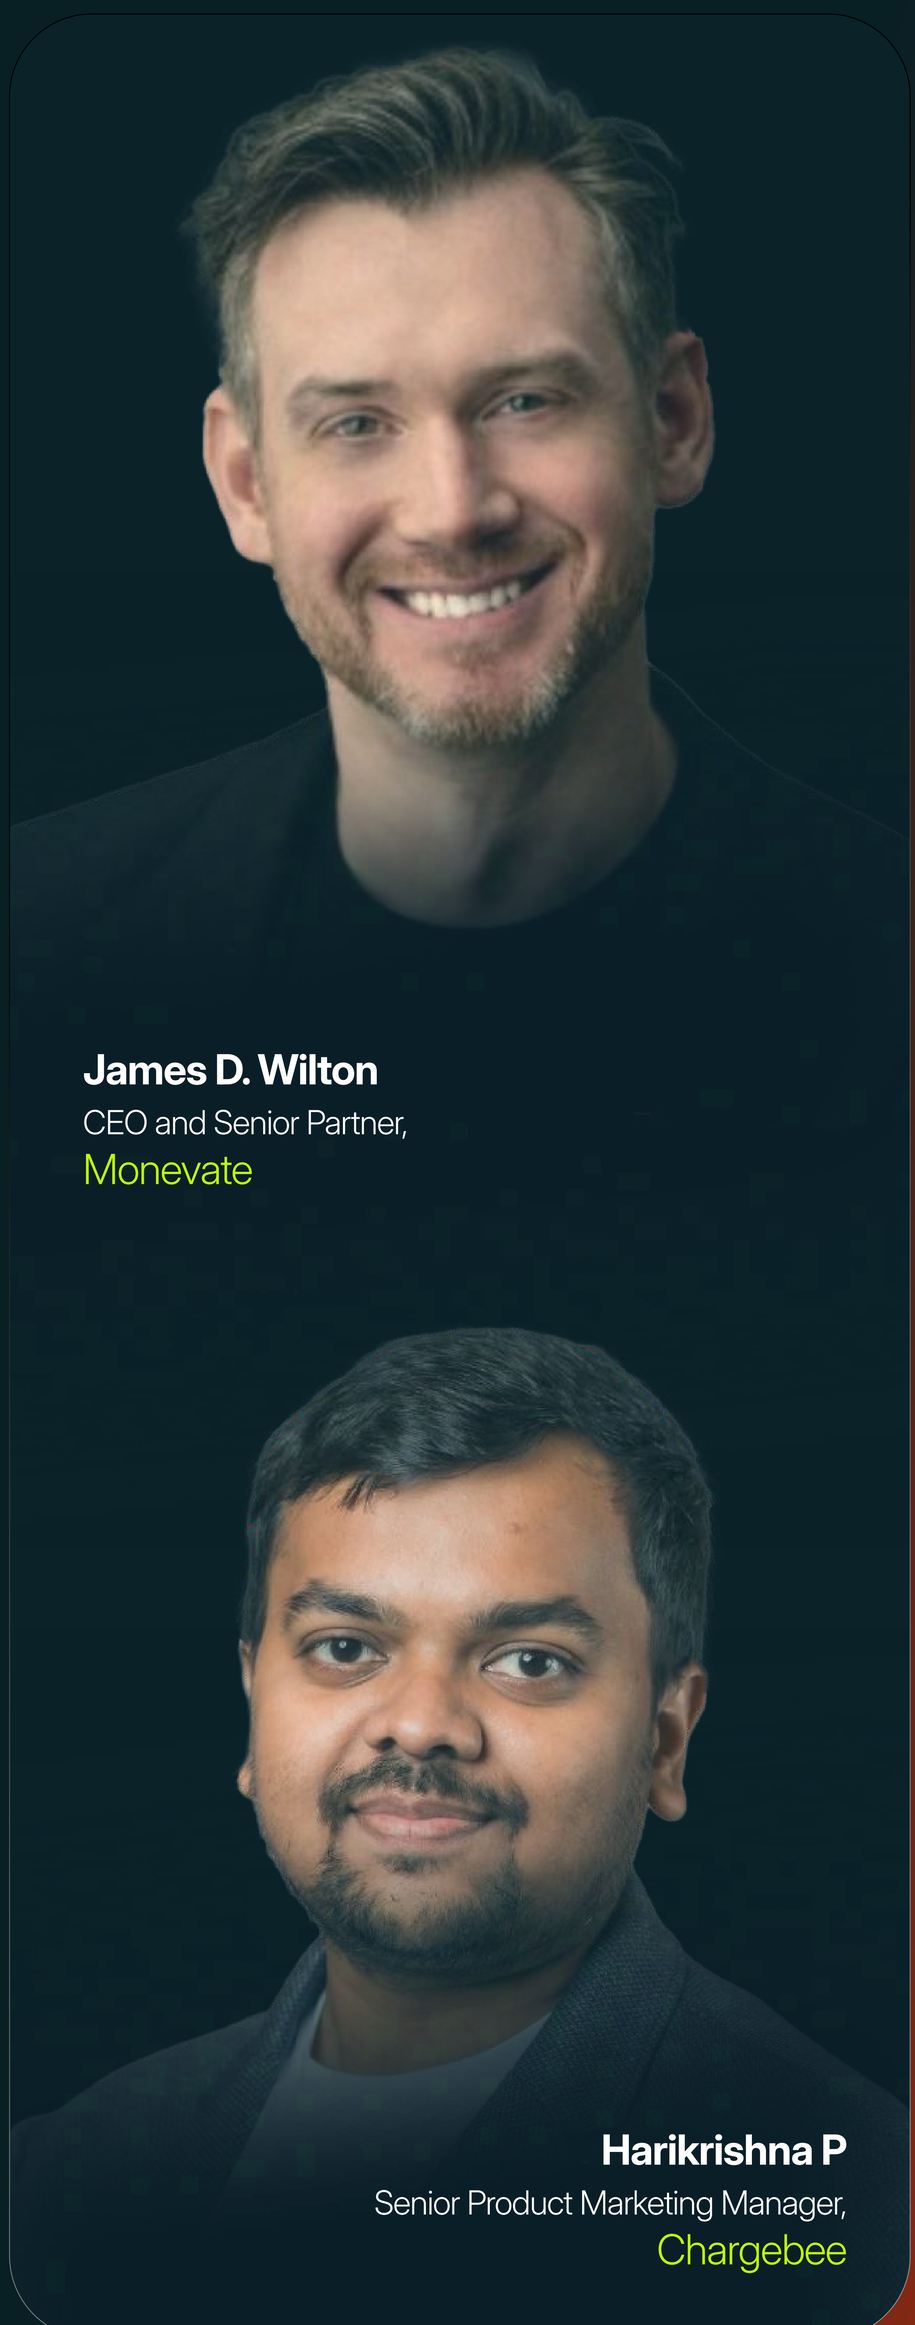

James D. Wilton
CEO and Senior Partner,
Monevate

Harikrishna P
Senior Product Marketing Manager,
Chargebee

About Monevate

Monevate is a strategic advisory firm specializing in monetization and pricing for SaaS and AI companies. With 50+ transformations delivered in just 3 years and a 95% client referral rate, Monevate helps teams move from value creation to value capture. Their expertise spans:

- Maximizing revenue with value-aligned metrics
- Driving expansion through scalable pricing architecture
- Protecting margins with disciplined execution
- Creating new revenue streams by monetizing features, products, and AI

Learn More

About Chargebee

Chargebee is the leading revenue growth management platform, powering pricing innovation for 6,500+ subscription businesses. As your monetization and billing infra, Chargebee helps you:

- Operationalize flexible pricing: support subscriptions, usage-based, and hybrid models
- Experiment faster: launch and iterate on pricing without engineering bottlenecks
- Unify operations: keep Product, GTM, and Finance aligned with a single source of truth
- Scale reliably: ensure accurate billing, revenue recognition, and churn prevention

Learn More

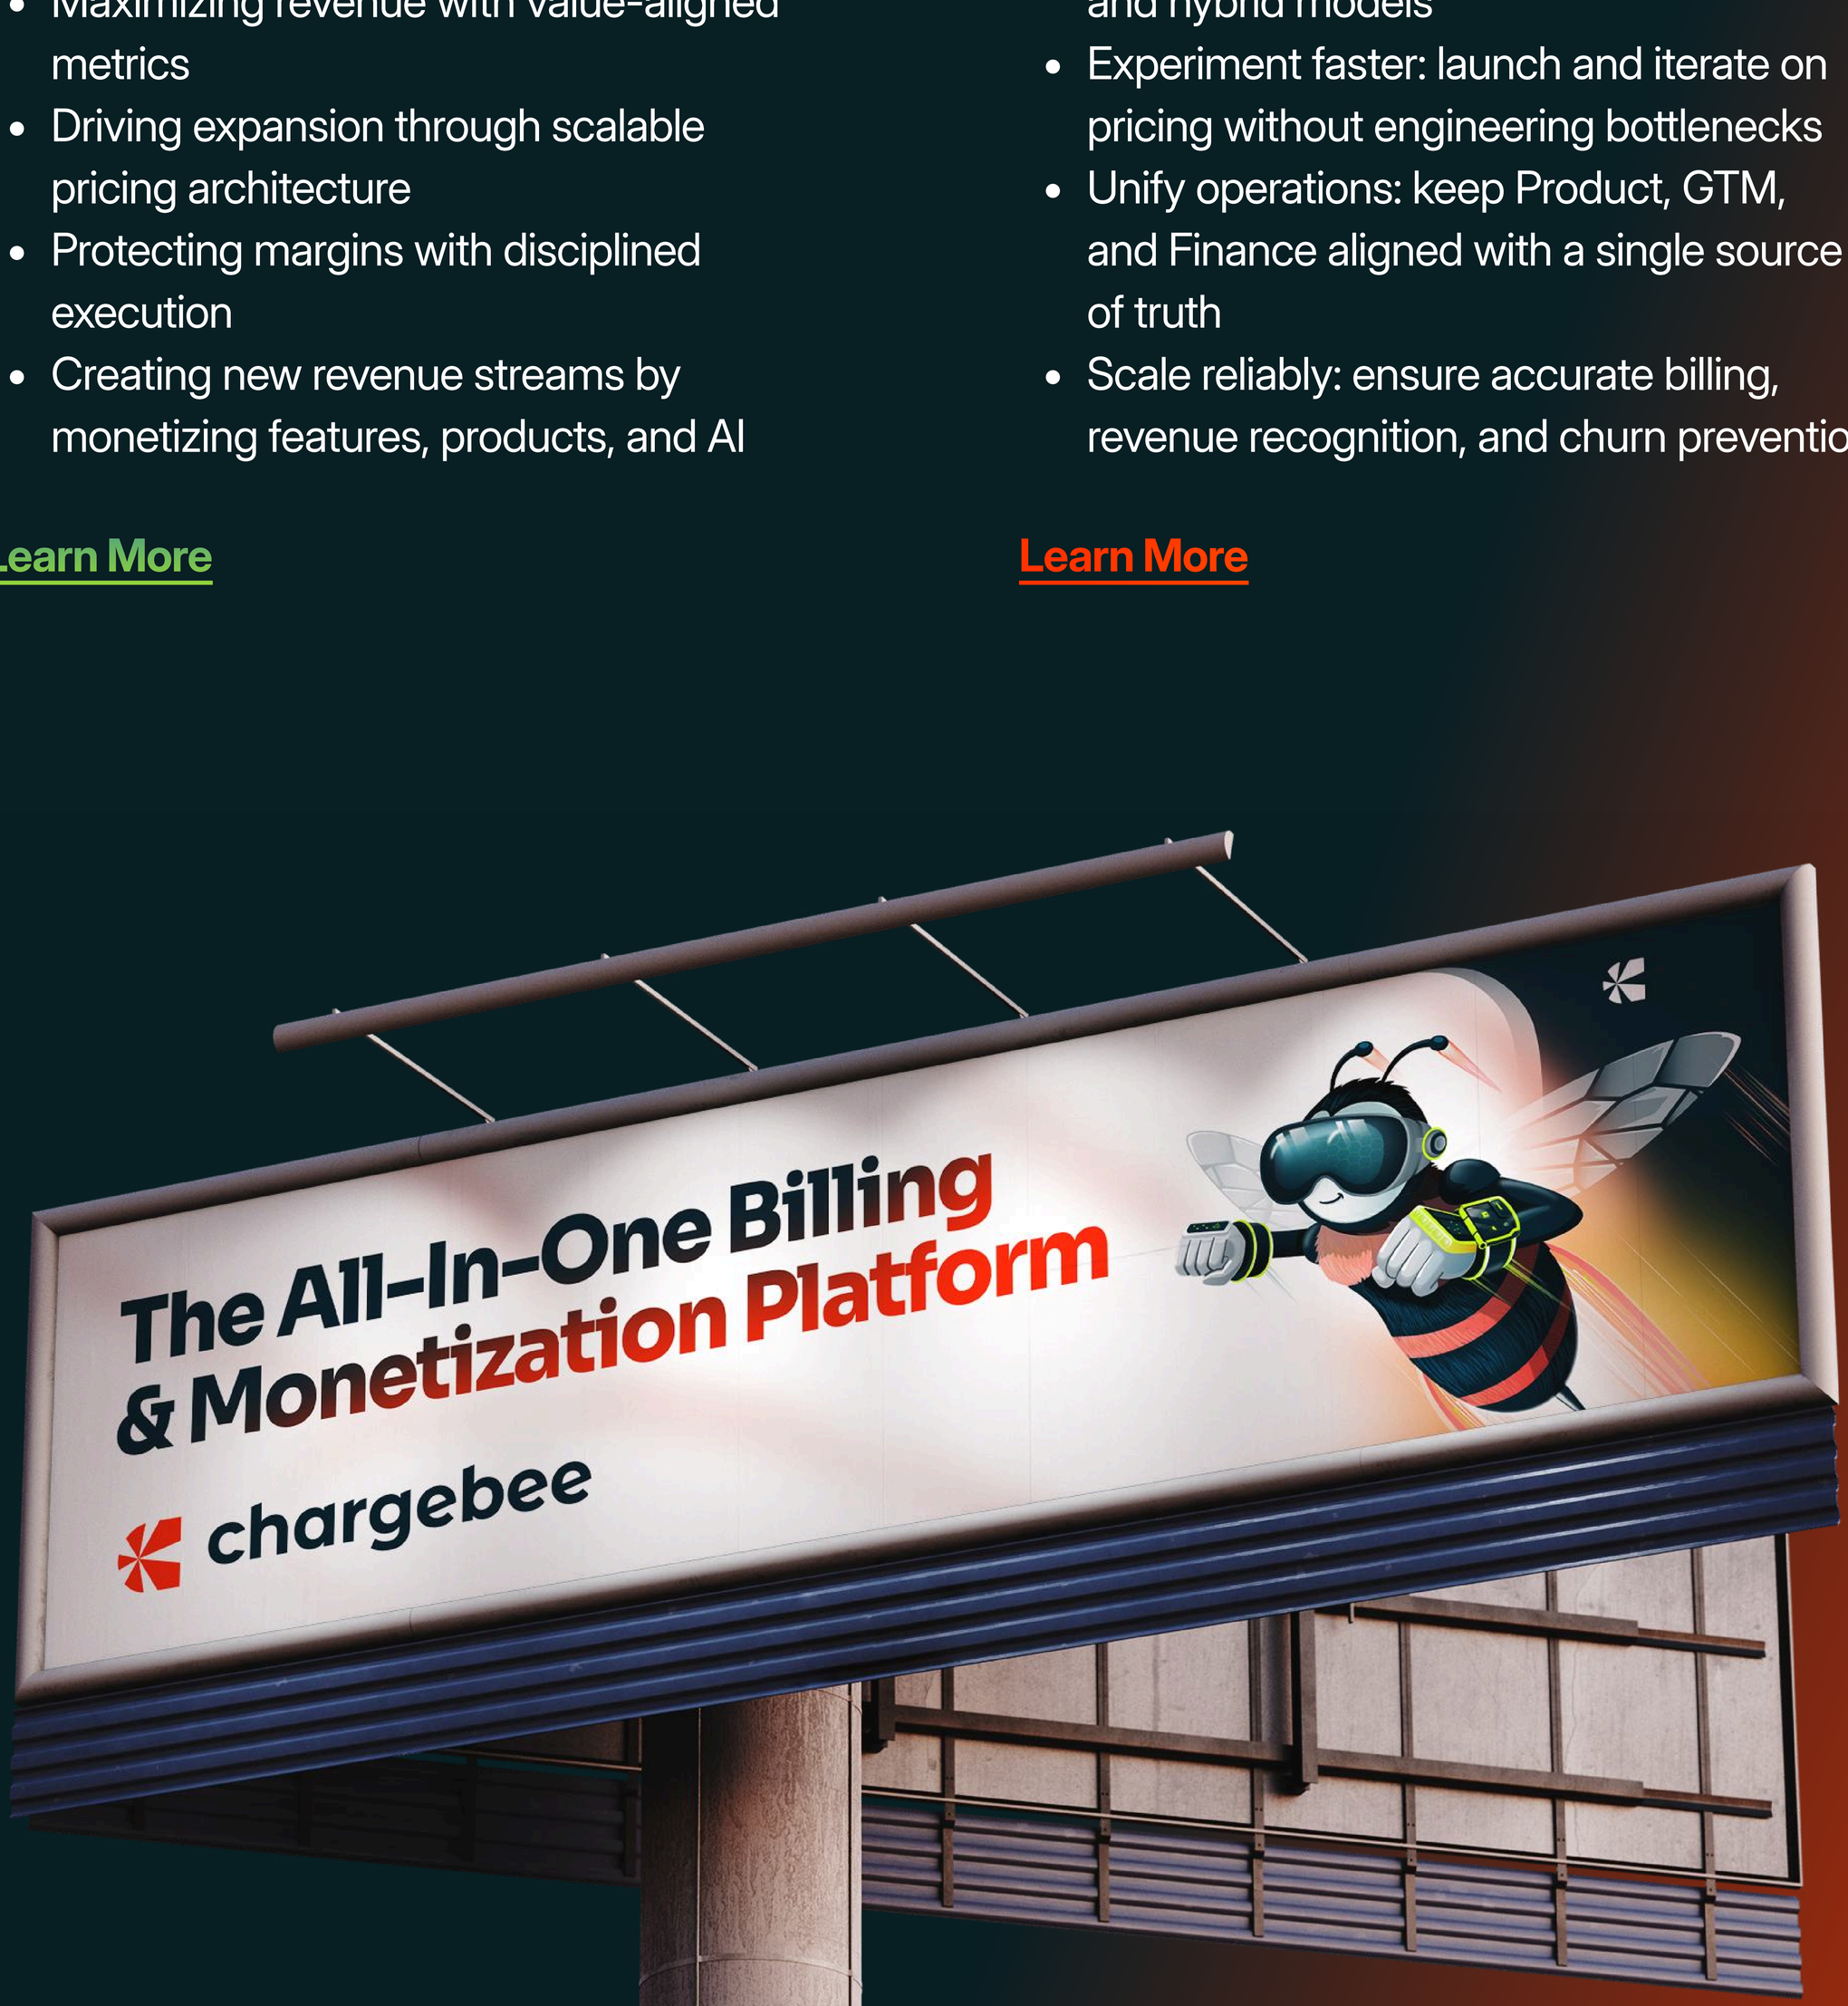

chargebee