PART 4

SAAS AND SUBSCRIPTION PRI
ARCHITECTURE

Metrics, Tiers, Freemium, and the Expansion Engine

SaaS pricing metric comparison (per-seat, usage-based, outcome-based, hybrid), per-seat ARR mechanics with minimum commitment and true-up structure, usage-based revenue model and committed vs. consumed ARR, freemium unit economics with effective CAC calculation, conversion gate design, tier architecture principles with pricing formula, annual vs. monthly discount sizing model, land-and-expand ARR model with NRR decomposition, lifetime cost of SaaS discounting, and the price increase execution framework with churn ROI model.

SECTION 1
SAAS PRICING FUNDAMENTALS

SaaS and Subscription Pricing Architecture

Software-as-a-Service pricing is the domain where pricing strategy has the highest financial leverage and the most direct impact on the metrics that determine company valuation. In SaaS, pricing is not simply a revenue decision — it is an architecture decision. The pricing model determines how revenue is recognized, how customers expand, what the churn dynamics look like, how the sales motion is structured, and ultimately what multiple the business commands from investors. A SaaS company with excellent product-market fit but poor pricing architecture will consistently underperform its potential. A company with mediocre product-market fit but excellent pricing architecture can sustain surprising financial performance while the product catches up.

The foundational choice in SaaS pricing is the pricing metric — the unit by which value is measured and charged. Per-seat (by user count), per-usage (by consumption), per-outcome (by value delivered), or per-feature-tier (by capability level) — each metric creates a different financial model with different expansion dynamics, different churn risk, different sales complexity, and different gross margin implications. Choosing the right pricing metric is the most consequential pricing decision a SaaS CFO makes, and it is almost never revisited with the rigor it deserves.

This part covers the complete financial architecture of SaaS pricing: per-seat vs. usage-based vs. outcome-based metrics, freemium economics and the conversion funnel, tier design and packaging, annual vs. monthly pricing incentives, expansion pricing and the land-and-expand model, discounting discipline and its effect on ARR metrics, and the price increase playbook — including how to execute a price increase that improves NRR without generating destructive churn.

1.1 The SaaS Pricing Metric Choice

Pricing Metric	Structure	Expansion Dynamic	Churn Risk	Best Fit
Per Seat (user)	Fixed price per user/month or year	Expands as headcount grows; organic upsell	High if users stop using; seat count shrinks	Collaboration tools, CRM, project management

Pricing Metric	Structure	Expansion Dynamic	Churn Risk	Best Fit
Usage-Based (consumption)	Pay per unit consumed (API calls, GB, MWh)	Expands naturally with customer success	Low; customers only pay for what they use	Infrastructure, data, AI inference, communications
Tiered Feature (good/better/best)	Fixed tiers at increasing price and capability	Upsell to higher tier as needs grow	Medium; value proposition per tier must be clear	Productivity tools, analytics, marketing platforms
Outcome-Based (value)	Price as % of value delivered or outcomes achieved	Expands as customer grows and value scales	Very low; customer wins when vendor wins	Revenue intelligence, collections, insurance-linked
Flat Rate (all-inclusive)	Single price for unlimited usage and users	No natural expansion; revenue flat per customer	Medium; simple but leaves money on table	Early stage; undifferentiated market; commodity
Hybrid (seat + usage)	Base seat license + usage-based overage	Dual expansion: users and consumption	Medium; complexity can confuse buyers	Mature SaaS with diverse customer profiles

SECTION 2
PER-SEAT PRICING ECONOMICS

Per-Seat Pricing: The Classic SaaS Model and Its Financial Architecture

Per-seat pricing — charging a fixed amount per user per month or year — remains the most widely used SaaS pricing model because of its simplicity, predictability, and alignment with how organizations think about software budgets. Finance teams understand 'X dollars per user per year.' Procurement teams can model total cost as headcount scales. Sales teams can compute deal size instantly. These qualities make per-seat pricing the lowest-friction model for enterprise sales and the most straightforward model for financial forecasting.

2.1 Per-Seat ARR Mechanics

PER-SEAT ARR BUILD AND EXPANSION

```
ARR = Active Seats x Annual Revenue Per Seat (ARPS)

New Business ARR = New Logos x Avg Initial Seat Count x ARPS
Expansion ARR = Existing Customers x Seat Growth Rate x ARPS
Contraction ARR = Existing Customers x Seat Reduction Rate x ARPS
Churn ARR = Churned Customers x Avg Seats at Churn x ARPS

Example: SaaS CRM tool
  Beginning ARR: $12,000,000  (200 customers x avg 30 seats x $2,000/seat/yr)
  New Business:  +$2,400,000  (40 new customers x 25 seats x $2,400/seat/yr)
  Seat Expansion: +$1,200,000 (existing customers add avg 5 seats)
  Seat Contraction: -$360,000 (some customers reduce seats at renewal)
  Churn:         -$840,000   (7 customers lost at avg $120K ARR each)
  Ending ARR:    $14,400,000

Net New ARR:     +$2,400,000  (20% growth)
NRR: ($12M + $1.2M - $0.36M - $0.84M) / $12M = 100% exactly
```

2.2 Per-Seat Pricing Risks and Mitigations

Per-seat pricing has two structural weaknesses that the CFO must manage. First, customers have an incentive to minimize seat count — to share logins, to not activate seats for occasional users, and to reduce seats aggressively at renewal even if usage has been healthy. This 'seat compression' dynamic can cause ARR contraction even among satisfied customers if the renewal process is not actively managed. Second, per-seat pricing does not expand with the customer's success or value realization — a customer who generates $10M in ROI from the product pays the same per-seat price as a customer who generates $100K in ROI, which means high-value customers are systematically undercharged.

The primary mitigation for seat compression is minimum commitment structures: requiring customers to commit to a minimum seat count for the contract term, with true-up provisions that charge for actual usage above the minimum but provide no credit for usage below. This converts the per-seat model into a hybrid committed/actual model that protects ARR floor while allowing upside from expansion.

MINIMUM COMMITMENT + TRUE-UP STRUCTURE

```
Customer commits to 50 seats minimum at $2,000/seat/yr = $100,000 ARR floor
Actual usage tracked monthly via admin dashboard

Scenario A: Customer uses 60 seats (above commitment)
  True-up: 10 additional seats x $2,000 = $20,000 billed at year-end
  Total: $120,000  |  ARR protected and expanded

Scenario B: Customer uses 40 seats (below commitment)
  No credit for unused seats; minimum commitment stands
  Total: $100,000  |  ARR protected despite usage decline

Renewal strategy: If customer used only 40 seats, renewal discussion
  centers on the 10 underutilized seats — adoption problem, not pricing problem.
  Solving adoption is more valuable than reducing the commitment.
```

SECTION 3

USAGE-BASED PRICING ECONOMICS

Usage-Based Pricing: Aligning Revenue with Customer Value

Usage-based pricing (UBP) — charging customers based on how much they consume of the product rather than how many users have access — is the fastest-growing pricing model in enterprise SaaS. Companies like Snowflake, Twilio, Stripe, AWS, DataStax, and OpenAI have built large businesses on usage-based models. The financial logic is compelling: revenue expands naturally as customers succeed and use more, creating a growth flywheel where the best customers also become the largest revenue contributors without requiring a separate upsell motion.

3.1 Usage-Based Revenue Model

Usage-based revenue has a fundamentally different financial profile than seat-based revenue. Revenue is variable rather than fixed — it fluctuates with customer usage patterns, which can be volatile, seasonal, or subject to rapid expansion and contraction. This variability creates challenges for financial forecasting (it is

harder to predict committed revenue) and for the income statement (revenue recognition is more complex because the performance obligation is satisfied over time as usage occurs, and the variable consideration must be estimated).

USAGE-BASED REVENUE MODEL

```
Monthly Revenue = Sum over all customers of (Units Consumed x Price per Unit)

Example: API platform at $0.008 per API call
  Customer A: 12M calls/month x $0.008 = $96,000/month
  Customer B: 3.2M calls/month x $0.008 = $25,600/month
  Customer C: 450K calls/month x $0.008  = $3,600/month
  Total MRR (100 customers):              $580,000/month

Revenue Expansion from Customer Success:
  Customer A grew from 8M to 12M calls (+50% in 6 months)
  Revenue grew from $64K to $96K with zero additional sales effort
  Net Revenue Retention calculation:
    If all customers' usage grew 30% on average: NRR = 130%
    If all customers' usage flat with 5% customer churn: NRR = 95%
    Usage growth is the primary NRR driver in UBP models
```

3.2 Committed ARR vs. Consumed ARR

A critical distinction in usage-based SaaS finance is between committed ARR (the contracted minimum the customer has agreed to pay) and consumed ARR (the actual usage-based revenue in the period). This distinction matters for several reasons: revenue forecasting (committed ARR is predictable; consumed ARR above commitment is variable), balance sheet treatment (committed ARR creates contract liabilities for pre-paid minimums; consumed ARR is recognized as usage occurs), and investor communication (many growth-stage UBP companies present both metrics to show the floor versus the run-rate upside).

CFO INSIGHT

The NRR metric is meaningfully different between seat-based and usage-based SaaS companies, and investors understand this. A usage-based SaaS company with 130% NRR is achieving it through organic usage growth — customers using more of the product as their own businesses grow. A seat-based company with 130% NRR is achieving it through active upsell and cross-sell motions that require sales capacity. The underlying quality of the NRR is different: usage-based NRR has lower CAC per expansion dollar because expansion happens naturally. When presenting NRR to a board or investor, show the decomposition between usage expansion (organic) and new seat/feature upsells (active sales). Organic expansion NRR commands a higher valuation multiple than active-sales NRR.

SECTION 4

FREEMIUM ECONOMICS: THE PAID CONVERSION FUNNEL

Freemium: The Financial Logic of the Free Tier

Freemium — offering a permanently free version of the product to attract users who may eventually convert to a paid tier — is one of the most debated pricing strategies in SaaS. Its proponents argue that it is the most efficient customer acquisition channel, reducing CAC by allowing customers to self-educate and self-qualify before engaging with a sales team. Its critics argue that it creates a massive cost liability (serving free users is not free), attracts users who will never convert, and trains customers to expect full functionality without payment. Both sides are right in different contexts. The CFO must model the economics of freemium rigorously before launching or expanding a free tier.

4.1 The Freemium Financial Model

FREEMIUM UNIT ECONOMICS

```
Freemium Funnel:
  Monthly Free Sign-Ups:          10,000
  Free-to-Paid Conversion Rate:      4.2%
  Monthly New Paid Customers:          420

Cost of Free Tier:
  Avg Infrastructure Cost per Free User: $0.35/month
  Monthly Free User Base: 85,000 (avg; includes multi-month free users)
  Monthly Free Tier Cost: 85,000 x $0.35 = $29,750

Revenue from Converted Users:
  Monthly New MRR: 420 x $85 avg MRR = $35,700
  LTV of converted user: $85 x 70% GM / 3% monthly churn = $1,983

Effective CAC from Freemium Channel:
  Monthly cost: $29,750 infra + $15,000 free tier support = $44,750
  Monthly new paid customers: 420
  Effective CAC: $44,750 / 420 = $106.55 per paid customer

LTV:CAC Ratio: $1,983 / $107 = 18.5x  (excellent)
  Compare to paid acquisition CAC: typically $300–$800 in B2C SaaS
```

4.2 Designing the Free-to-Paid Conversion Gate

The most important design decision in freemium is the conversion gate — the point at which the free user encounters a limitation that can only be resolved by upgrading to a paid plan. The gate must be designed so that it is encountered by users who have already experienced enough value to be motivated to upgrade (not by new users who have not yet formed a habit), and so that the upgrade friction is low enough that motivated users actually convert rather than churning.

There are three primary gate architectures. The usage gate blocks specific features (export, integrations, advanced analytics) from free users but allows unlimited use of core features. The capacity gate allows free users to use all features but limits volume (5 projects, 1,000 contacts, 10GB storage). The time gate allows full access for a trial period and then restricts to a limited free tier. Each gate has different conversion dynamics and different implications for the user experience. Usage gates tend to convert users who need the specific blocked feature — conversion is intentional and high-quality. Capacity gates convert users who have grown beyond the free limits — conversion is driven by success, which is the ideal trigger.

Gate Type	Trigger	Conversion Quality	Risk	Best For
Feature Gate	User needs a specific paid-only feature	High — user knows exactly what they want	Gate on wrong features; low trigger rate	Collaboration tools; analytics; integrations
Capacity Gate	User exceeds usage limit (contacts, projects, storage)	Very High — user has proven value before converting	Setting limits too high delays conversion	CRM, storage, project management
Time Gate (trial)	Free trial period expires	Variable — depends on engagement during trial	Low-engagement users churn; don't convert	Complex products needing demonstration period
Team/Org Gate	User wants to invite colleagues	High — social network drives conversion	Requires viral/collaboration product mechanic	Collaboration tools, communication platforms

SECTION 5
TIER DESIGN AND PACKAGING

Tier Design: Good, Better, Best — and Why Most Companies Get It Wrong

Most SaaS products are sold through a tiered pricing structure — commonly labeled Starter/Professional/Enterprise, Basic/Pro/Business, or some variation of Good/Better/Best. The logic is appealing: offer different capability levels at different prices to capture different WTP segments. The execution is frequently poor: tiers that are too similar (no meaningful differentiation between Starter and Professional), too dissimilar (enormous price jumps with no intermediate option), or designed around features the product team built rather than the value outcomes customers actually care about.

5.1 Tier Architecture Principles

Effective tier design follows three principles. First, each tier must be anchored to a distinct customer segment with distinct value drivers — not just a different feature list. The Starter tier serves a specific persona with specific needs; the Professional tier serves a meaningfully different persona who needs something the Starter tier genuinely cannot provide. Second, the price gaps between tiers must reflect the

value differential between them, not an arbitrary multiple. If Professional is twice the price of Starter, it should deliver twice the value — or the Professional tier will struggle to sell. Third, the highest-priced tier should contain the features that create the most lock-in and the deepest integration — the features that make it most costly for the customer to switch to a competitor.

TIER PRICING DESIGN FRAMEWORK

```
Step 1: Define segment personas for each tier
  Starter: Individual contributor; needs core functionality; price-sensitive
  Professional: Small team (3–15 users); needs collaboration + reporting
  Enterprise: Large team (50+); needs security, SSO, compliance, SLA

Step 2: Identify value differentials between tiers
  Starter to Professional: Collaboration value = ~$180/user/yr (EVC-based)
  Professional to Enterprise: Security/compliance value = ~$600/user/yr

Step 3: Set prices to capture ~50%–65% of value differential
  Starter: $120/user/yr  (baseline; covers variable cost plus margin)
  Professional: $240/user/yr  (captures ~67% of collaboration value above Starter)
  Enterprise: $600/user/yr  (captures ~73% of security/compliance above Pro)

Step 4: Design feature gates to enforce tier boundaries
  Features that generate collaboration value -> Professional only
  Features that generate compliance/security value -> Enterprise only
  Core functionality that proves base value -> all tiers including free
```

5.2 Packaging vs. Pricing: The Add-On Architecture

An alternative to pure tiered pricing is a modular packaging architecture: a base platform at a core price plus optional add-on modules that customers purchase based on specific needs. Add-on pricing captures value from the specific modules a customer needs without requiring them to move to an entirely new tier. It also enables expansion within existing customers without a full-tier upsell — a customer on the Professional tier who needs one specific Enterprise feature can purchase that feature as an add-on rather than migrating their entire team to Enterprise pricing.

The financial advantage of add-on architectures is that they create multiple expansion paths, each with its own upsell conversation and its own budget approval process. Instead of one upsell (tier migration) that requires a large budget discussion, the account team has four or five smaller upsell conversations (add-on purchases) that can each be approved at a lower budget authority level. This dramatically increases the

expansion ARR velocity in existing accounts.

SECTION 6
ANNUAL VS. MONTHLY PRICING AND CONTRACT TERMS

Annual vs. Monthly: The Cash Flow Architecture of Subscription Pricing

The choice between annual and monthly billing — and how to incentivize one over the other — is a pricing decision with significant financial consequences for cash flow, churn dynamics, and the working capital structure of a SaaS business. Annual contracts paid upfront provide a substantial cash flow benefit (the company receives 12 months of cash at contract signing), reduce churn (customers who have paid annually are less likely to cancel in month 4 than customers who are billed monthly), and simplify revenue forecasting. Monthly billing preserves customer optionality (which reduces friction in the purchase decision) and may increase overall market penetration by eliminating the commitment barrier.

6.1 The Financial Case for Annual Contracts

ANNUAL VS. MONTHLY FINANCIAL COMPARISON

```
Customer A: Monthly billing at $400/month
  Monthly revenue: $400 | Monthly CAC recovered: $400 each month
  At 3% monthly churn: 50% probability of still being a customer at month 12
  Expected LTV (12 months): $400 x 8.7 avg months = $3,480
  Cash received by month 6: $2,400 (6 payments)

Customer B: Annual billing at $3,840/year ($320/mo effective; 20% discount)
  Cash at contract signing: $3,840 (full year upfront)
  Annual churn risk: 12% (vs. 3% monthly = 31% annual effective for monthly)
  Cash received by month 6: $3,840 (all at month 1)

Working Capital Advantage of Annual vs. Monthly:
  At 1,000 customers on $400/month plans:
    Monthly billing: ~$400K MRR; no upfront cash beyond monthly billing
    Annual billing: $3.84M upfront per cohort; $3.84M in deferred revenue
  Annual billing provides ~$3.44M more cash in Year 1 vs. monthly
  This cash funds growth without external financing
```

6.2 Structuring the Annual Discount

Most SaaS companies offer a discount for annual prepayment — typically 15% to 20% off the monthly equivalent price. The CFO must size this discount to ensure it is financially positive: the company should be willing to give up at most the present value of the monthly payment stream, adjusted for the churn reduction benefit of having a committed customer.

MAXIMUM RATIONAL ANNUAL DISCOUNT CALCULATION

```
Monthly Price: $400  |  Annual equivalent (no discount): $4,800

Monthly billing economics (probabilistic):
  Monthly churn rate: 2.5%
  Expected months before churn: 1/2.5% = 40 months
  Expected annual revenue from monthly customer: $400 x min(12, 40) = $4,800
  But: 26% probability of churning within year 1 (1-(1-0.025)^12)
  Expected annual revenue (risk-adjusted): $4,800 x 74% = $3,552

Annual billing eliminates churn-within-year risk:
  Guaranteed annual revenue: $4,800 x (1 - annual discount %)
  Set guaranteed annual >= expected monthly annual revenue:
    $4,800 x (1 - d) >= $3,552
    1 - d >= 0.74
    d <= 26%  (maximum rational discount is 26% in this scenario)

Standard industry discount: 15%–20%  (conservative; leaves margin of safety)
  At 20% discount: Annual price = $3,840; cost of discount = $960/customer
  Benefit: Guaranteed $3,840 vs. probabilistic $3,552 + better renewal rate
```

SECTION 7

EXPANSION PRICING AND THE LAND-AND-EXPAND MODEL

Expansion Pricing: Growing Revenue Inside Existing Accounts

Expansion revenue — additional revenue generated from existing customers through upsell (moving to a higher tier), cross-sell (purchasing additional products), or usage growth (consuming more of a usage-based product) — is the most valuable revenue in a SaaS business. It costs far less to generate than new logo revenue (no CAC for the expansion, or a fraction of new logo CAC), it arrives with pre-existing trust and integration, and it is a direct reflection of the customer's satisfaction and success with the product. The Net Revenue Retention (NRR) metric — which measures the total revenue retained and expanded from the existing customer base — is the single number that best captures the health of the expansion pricing engine.

7.1 The Land-and-Expand Financial Model

LAND-AND-EXPAND ARR MODEL

```
Year 1 — Land: Close 80 new customers at avg $18,000 ACV
  New ARR: $1,440,000  |  CAC: $12,000/customer -> Total Sales Cost: $960,000

Year 2 — Expand within cohort:
  Churn: 8% of Year 1 customers -> 74 surviving customers
  Expansion: 40% of survivors add avg $8,000 ARR (upsell/more seats)
  Expansion ARR: 30 customers x $8,000 = $240,000
  Expansion CAC: $2,000/customer (customer success motion; no new sales)
  Total cohort ARR Year 2: (74 x $18,000) + $240,000 = $1,572,000
  NRR: $1,572,000 / $1,440,000 = 109.2%

Year 3 — Continued expansion:
  Churn from Year 2 base: 7% -> 68 surviving original customers
  Additional expansion: another 35% expand
  Total cohort ARR Year 3: ~$1,680,000
  Cohort NRR Year 3: $1,680,000 / $1,440,000 = 116.7%

Key insight: Cohort NRR > 100% means cohort pays back CAC multiple times over
```

7.2 Pricing for Expansion: Designing the Upsell Path

Expansion revenue does not happen automatically — it requires a product and pricing architecture that creates natural expansion pressure at the right points in the customer journey. The most effective expansion triggers are those that arise organically from the customer's success: they have used up their capacity, they have grown their team, they have encountered a use case that requires a higher-tier feature. Forced upsells — pushing customers to upgrade before they have experienced the limitation — generate resistance and reduce customer satisfaction without the same revenue quality.

The CFO must model the expansion revenue contribution to NRR separately from the base retention. This decomposition reveals whether NRR is being driven by genuine customer success and value expansion or by aggressive upsell tactics that mask underlying retention weakness. A company with 110% NRR driven by 95% gross retention and 15% expansion is in a different financial position than one with 110% NRR driven by 85% gross retention and 25% expansion — even though the headline NRR is identical.

SECTION 8
DISCOUNTING DISCIPLINE IN SAAS

Discounting Discipline: The Hidden Destroyer of SaaS Economics

Discounting in SaaS is financially more damaging than in most other business models because of the compounding nature of recurring revenue. In a one-time sale, a 20% discount reduces the revenue from that sale by 20%. In a SaaS subscription, a 20% discount at contract signing reduces every renewal, every upsell, and every expansion conversation for the life of that customer — because the discounted price becomes the baseline for all future pricing conversations. A customer signed at $40,000 ARR who should have been at $50,000 ARR represents not just a $10,000 shortfall in Year 1 but potentially $30,000 to $50,000 in foregone ARR over a typical 4- to 6-year customer lifetime.

8.1 The Compounding Cost of a SaaS Discount

LIFETIME COST OF A 20% SAAS DISCOUNT

```
Deal A: Correct price $50,000 ARR; 5-year relationship; 10% annual expansion
  Year 1: $50,000  |  Year 2: $55,000  |  Year 3: $60,500
  Year 4: $66,550  |  Year 5: $73,205
  5-Year Total: $305,255

Deal B: 20% discount at signing; $40,000 ARR; same expansion rate
  Year 1: $40,000  |  Year 2: $44,000  |  Year 3: $48,400
  Year 4: $53,240  |  Year 5: $58,564
  5-Year Total: $244,204

Lifetime Cost of the 20% Discount: $305,255 - $244,204 = $61,051
  The $10,000 Year 1 discount compounds to $61,051 over 5 years
  At 75% gross margin: Lost contribution = $61,051 x 75% = $45,788

Plus: Discounted customer sets a pricing precedent that makes
  all future price increase conversations start from a lower base
```

CFO INSIGHT

The most dangerous SaaS discount is the one given to close the deal at end of quarter. Sales teams under quota pressure will offer 15% to 25% discounts in the final week of a quarter to meet their targets — discounts that would never be offered if the deal had been worked properly through the quarter. The CFO must build a reporting system that tracks the average discount rate by sales rep, by quarter, and by deal size. End-of-quarter discount rates that are materially higher than mid-quarter rates are a clear signal that the sales team is using price to compensate for insufficient pipeline development. The solution is pipeline management, not discount policies — but the data to identify the problem comes from rigorous discount tracking.

SECTION 9

THE SAAS PRICE INCREASE PLAYBOOK

SaaS Price Increases: Improving NRR Without Destroying Churn

The price increase is the highest-leverage improvement action available in a mature SaaS business — and the one executed most poorly. Many SaaS companies avoid price increases for years out of fear of customer backlash, even as inflation erodes their margins and competitors raise prices. When they finally act, they execute poorly: insufficient notice, inconsistent communication, no value reaffirmation, and no exceptions process for at-risk customers. The result is unnecessary churn from customers who would have accepted the price increase with better execution.

9.1 The Price Increase Execution Framework

Phase	Timing	Activities	CFO Role
Analysis	12+ weeks before	Model churn impact; identify at-risk accounts; design exception policy; set increase amount	Own the financial model; set acceptable churn threshold
Preparation	8–12 weeks before	Draft communication; train CS and sales; build value reaffirmation materials; prepare FAQs	Review communication for financial accuracy; approve exception criteria

Phase	Timing	Activities	CFO Role
Communication	6–8 weeks before	Send formal price increase notice to all affected customers; proactive outreach to top accounts	CFO-to-CFO outreach for strategic accounts above threshold
Retention	4–6 weeks before	CS team works at-risk accounts; exception approvals for accounts meeting criteria; retention offers	Approve exceptions; monitor early churn signals vs. model
Execution	At renewal date	New price takes effect; exceptions honored per approved list; track churn vs. model	Weekly monitoring vs. forecast; adjust exceptions if overrun
Review	60–90 days after	Analyze actual churn vs. model; customer feedback analysis; pricing team debrief	Financial impact vs. projected; lessons for next increase

9.2 Sizing the Price Increase and the Churn Tradeoff

PRICE INCREASE ROI MODEL

```
Current ARR Base: $25,000,000  |  Proposed Increase: 12%
Incremental ARR if 0% churn: $25M x 12% = $3,000,000

Churn Scenarios:
  0% incremental churn: Net NRR impact = +12%  |  Incremental ARR = $3,000,000
  3% incremental churn: Lost ARR = $25M x 3% = $750,000
    Net incremental ARR: $3,000,000 - $750,000 = $2,250,000  -> still positive
  8% incremental churn: Lost ARR = $25M x 8% = $2,000,000
    Net incremental ARR: $3,000,000 - $2,000,000 = $1,000,000 -> positive but thin
  12% incremental churn: Lost ARR = $25M x 12% = $3,000,000
    Net incremental ARR: $0  -> break-even; price increase NPV-neutral

Decision: Is 12% incremental churn likely?
  If price elasticity is low (SaaS with high switching costs): <3% likely
  If competitive alternatives are strong: 6%–10% possible
  Monitor early signals: inbound cancelation requests within 2 weeks of notice
```

9.3 Metrics and CFO Checklist

Metric	Formula / Definition	Benchmark
Net Revenue Retention (NRR)	(Beg ARR + Expansion - Contraction - Churn) / Beg ARR	>120% exceptional; >110% strong; 100% break-even
ARR per Seat (ARPS)	Total ARR / Total Active Seats	Track trend; declining = pricing erosion or tier migration down
Expansion ARR %	Expansion ARR / Total ARR	>20% is healthy; rising = land-and-expand working
Average Discount Rate	Total Discounts / Gross ARR Quoted	<10% target; >20% = sales discounting culture
Annual vs. Monthly Mix	Annual contract ARR / Total ARR	>60% annual preferred for cash flow and churn
Freemium Conversion Rate	Free-to-paid conversions / Total free signups	>3% healthy B2C; >8% strong product-led growth
Price Increase Churn Attribution	Churn directly attributed to price increase / Total churn	Track separately; should be <25% of total churn
Committed vs. Consumed ARR	Committed minimum / Total contracted ARR	Higher committed % = more predictable revenue base

- SaaS pricing metric reviewed annually: per-seat, usage-based, outcome-based, or hybrid evaluated against WTP research and competitive dynamics; pricing metric change modeled before execution.
- Annual contract discount policy enforced: discount calculated using maximum rational discount model; discounts above 20% require CFO approval with documented rationale.
- Expansion ARR reported separately from new logo ARR: NRR decomposed into gross retention, seat expansion, and tier upsell components; trend tracked monthly.
- Price increase analysis completed at least annually: churn impact modeled; at-risk account list developed; exception policy approved before communication begins.
- Discounting tracked by rep and by quarter: end-of-quarter discount rate compared to mid-quarter rate; systemic discounting addressed through pipeline management, not lower price floors.

Closing Perspective: SaaS Pricing as Financial Architecture

SaaS pricing is not a product decision that the CFO ratifies — it is a financial architecture decision that the CFO must lead. The choice of pricing metric determines the shape of the revenue growth curve. The tier design determines the expansion pathways that drive NRR. The annual vs. monthly policy determines the working capital structure and the churn dynamics. The discounting discipline determines whether the ARR metrics shown to investors actually represent the economics of the business. And the price increase playbook determines whether the company can maintain pricing power as it matures.

The CFO who owns all of these decisions — who can articulate the financial implications of a pricing metric change, model the NRR impact of a tier restructure, quantify the compounding cost of a 20% discount, and execute a price increase with precision — is the CFO who builds the recurring revenue foundation that sustains long-term enterprise value creation.

Part 5 covers Product and Physical Goods Pricing — MSRP, MAP, and channel pricing architecture, promotional pricing mechanics and the markdown optimization framework, keystone pricing, bundling and unbundling economics, international pricing and currency management, DTC vs. wholesale price architecture, and commodity pricing with hedging.

End of Part 4: SaaS and Subscription Pricing Architecture | Pricing Strategy — A 14-Part Series

eFuturesCFO | The Systems CFO Platform | efuturescfo.com